Organizational Culture, Adaptation and Performance

Francisco Brahm and Joaquin Poblete[1]

This version: June 2023

Abstract

Despite the agreement that culture affects organizational performance, we don't understand fully the mechanisms driving this relationship. Drawing from cultural evolution theory, we use a formal model to study a relatively understudied mechanism, namely, how organizational culture affects adaptation to a changing environment. Following prior literature, organizational culture is differentiated into two dimensions, namely the "beliefs" about the organization's function in the environment (i.e., "what problems we solve and how we solve them") and the "prosocial norms" about internal integration (i.e., "how do we come together as a unit"). Our main findings are: i) when beliefs are a more accurate representation of the organization's true function, the timing of innovation improves (i.e., innovation focusses on periods of environmental change) leading to better adaptation to changes; concurrently, the use of tradition increases, as well as the effort in diffusing it; these two effects entail a reduction in the tradition-innovation trade-off; ii) more intense prosocial norms improve performance via increased effort in diffusing tradition (creating a "cushion" to absorb environmental shocks), but if intensity is sufficiently high, tradition's crowding out of innovation and adaptation becomes too large, reducing performance—and thus, prosocial norms display an inverted-U relationship with performance; and iii) beliefs and prosocial norms are complements: knowing better when to innovate is more beneficial if the organization is better at sharing the innovations (and vice versa). In sum, culture can favor adaptation by improving the timing of innovation, reducing the tradition vs. innovation trade-off and serving as a buffer to shocks.

Keywords: Cognition, Cultural evolution, Learning, Organizational adaptation, Organizational culture, Shared beliefs, Prosocial norms.

Running header: Organizational culture and adaptation

1. Introduction

Over the last few decades, research on organizational culture—broadly defined as the beliefs, values and norms shared by organizational members—has been accumulating across different fields of academic inquiry (for organization theory, see Giorgi et al., 2015 and Weber and Dacin, 2011; for strategy, see Gibbons et al., 2021b; for organizational behavior, see Chatman and O'Reilly, 2016, Hartnell et al., 2011 and 2019, and Srivastava and Goldberg, 2017; for finance, see Grennan and Li,

---

[1] Brahm: London Business School, 26 Sussex Pl, London NW1 4SA, UK, fbrahm@london.edu (corresponding author); Poblete: School of Management, Pontificia Universidad Catolica de Chile (PUC), Av. Libertador Bdo. O'Higgins 440, 4to piso, Santiago 8331010, Santiago, Chile, joaco.poblete@gmail.com. This paper has benefited from the comments by Arianna Marchetti, Michael Jensen, Michael Jacobides, Christoph Loch, Freek Vermeulen and Robert Gibbons; by participants in seminars at the Michigan Ross School of Business, the London Business School and the School of Management at PUC; and attendees of the 2nd Berkeley Haas Culture Conference, 2021 TOM annual conference and 2022 AOM Annual Conference. An earlier version of his paper circulated as "The evolution of organizational culture."

2022 and Gorton et al., 2022; for economics, see Hermalin, 2012). For this literature, it is self-evident that organizational culture affects performance. Recent evidence from firm executives confirms this: 91% of 1,348 executives in North America agree that culture is important or very important for performance at their firm (Graham et al., 2022). However, several authors in this field suggest that, despite progress, there is much work to be done to understand in detail the mechanisms through which culture affects performance (O'Reilly and Chatman, 2016; Chatman and Srivastava, 2021; Chatman et al., 2014; Gibbons et al., 2021b; Grennan and Li, 2022).

In this paper we study how organizational culture affects a relatively understudied channel, namely, how it affects the capacity of the organization to adapt to environmental changes. Following Schein (2010) and others[2] we conceptualize organizational culture as possessing two broad "dimensions": *beliefs* about the problem the firm solves (i.e., its function) and the associated mean-ends relations (i.e., if we do "A", we obtain "B"); and *pro-social norms*, that is, cooperative norms that prioritize help to colleagues and contributions to the whole organization (rather than focusing on one's own payoff). We study how each of these two dimensions of culture affects adaptation by using a formal model drawn from cultural evolution theory (Boyd and Richerson, 1985 and 2005; Henrich, 2016 and 2020; Cavalli-Sforza and Feldman, 1981; Boyd, 2018; Mesoudi, 2017; Brahm and Poblete, 2021 and 2022).

Regarding beliefs, we draw from research on organizational cognition (Tripsas and Gavetti, 2000; Tripsas, 2009; Eggers and Kaplan, 2009 and 2013; Helfat and Peteraf, 2015; Barr et al., 1992) that shows that beliefs, as a representation of underlying reality (Puranam, 2018; March and Simon, 1958; March, 1991), can have a different level of accuracy or verisimilitude (ibid). We find that belief accuracy allows the organization to better interpret signals from its interaction with the environment, and thus, decide better when it is necessary to innovate, improving the organization's adaptation to environmental changes. Interestingly, we also find that more accurate beliefs, due to better timing of innovation, increases the frequency of tradition; in turn, this promotes an expansion of cooperative

[2] Kocak and Puranam (2022) distinguishes between causal codes and moral codes in culture and DiMaggio (1994) distinguishes between constitutive beliefs (taken-for-granted cognitive categories and schema necessary for parties to think) and regulative beliefs (norms, values, and conventions that may reshape an individual's pursuit of self-interest).

effort geared towards making tradition more efficient. The combination of these two effects means that more accurate beliefs reduce the trade-off that arises between tradition and innovation (March, 1991; O'Reilly and Tushman, 2013; Benner and Tushman, 2006) because it helps both perform better.

Regarding prosocial norms, we find that more intense prosocial norms make the organization's tradition—its accumulated practices, skills, technology—more efficient as agents exerts in effort in transmission; this improvement in performance, serves to buffer the organization from external shocks. However, by promoting tradition, the intensity of prosocial norms also reduces the capacity of the organization to innovate and track the environment, becoming increasingly liable to environmental changes; if intensity is sufficiently high, this effect overwhelms the efficiency gains, and performance suffers. Thus, against received wisdom regarding pro-sociality (Organ et al, 2005; Fehr, 2018; Rand and Nowak, 2013; Liao et al, 2022) the intensity of pro-social norms displays an inverted-U relation with organizational performance.

The formal model also shows that belief accuracy and intensity of prosocial norms are complements: the impact of each cultural dimension on performance is more positive (or less negative) if the other one increases. While not obvious a priori, the model provides a clear intuition: well-timed innovation produced by accurate beliefs produces more benefits when attached to a more efficient internal transmission of the innovation (driven by more intense prosocial norms)[3]. We are not aware of prior research exploring the interaction between these two dimensions of organizational culture.

In sum, we show that by shifting beliefs' accuracy and prosocial norms' intensity, two building blocks of any organizational culture, the adaptability of the organization to changes is affected. Organizational culture can improve the timing of innovation, make tradition more efficient and thus reduce the tradition vs. innovation trade-off, and serve, up to a point, as a buffer from shocks.

[3] This also operates in the reverse direction: efficient transmission of tradition driven by pro-sociality is more valuable when that tradition can be timely updated when changes occur (by way of accurate beliefs).

1.1. Contributions to the literature

We offer several contributions to the organizational culture literature. On beliefs, the contribution is that i) cultural strength doesn't need to come at the expense of adaptation (cf., Corritore et al, 2020), and ii) we offer a way to break the trade-off between efficiency and adaptation attached to cultural strength that addresses some limitations of prior attempts (Chatman et al, 2014; Gibbons et al, 2021c).

Let's unpack i). The literature in organizational culture has shown, via theoretical models (Cremer, 1993; Van den Steen, 2010a; Gibbons et al., 2021a; Fehr, 2018) and robust empirical evidence (Marcheti, 2019; Marcheti and Puranam, 2020; Grennan, 2020; Grennan, 2022; Graham et al., 2022; Sørensen, 2002; Corritore et al., 2020; Chatterji et al., 2016; Weber and Camerer, 2011), that high consensus/intensity in beliefs (i.e., a strong culture) has a positive impact on the alignment of expectations, the smoothness of communication and coordination, and the overall efficiency of the organization. However, this generates a trade-off: a strong culture, by limiting diversity and thus the exploration of new/innovative ideas, constrains the capacity to adapt to environmental changes (March, 1991; Sorensen, 2002; Van den Steen, 2010b; Corritore et al., 2020; Hermalin, 2012). Corritore et al. (2020) provides compelling, large-scale evidence supporting this trade-off. In our paper, we show that even if beliefs are homogeneous—in our model, all organizational members have the same beliefs—adaptation improves when their accuracy improves. Thus, we provide a mechanism that can compensate for the loss in adaptability from belief homogeneity, and thus minimize, or even break, the trade-off.

Regarding ii), we join two previous efforts, Chatman et al (2014) and Gibbons et al (2021c), that have tried to break away from the "cultural strength trade-off". Chatman et al. (2014) argues that a strong culture may favor adaptation by strongly emphasizing adaptability in its content. A limitation of assuming adaptability in the culture's content is that adaptability is relevant in some settings, such as high-tech, and not others, such as stable and mature industries. Our model doesn't prescribe a specific content for beliefs, and thus we don't "assume" away the trade-off.

Regarding the second effort, Gibbons et al. (2021c) draws from Kreps (1990)—who interpreted a repeated-game equilibrium as a shared understanding of "how things are done," which usually takes the form of general principles/beliefs about proper behavior (e.g., “when in conflict, always defer decision making to the senior party,” “seek fairness within the organization,” “prioritize customers”—to argue that when the circumstances change (e.g., a change in payoffs of the game) and agents must adapt in response, the presence of principles/beliefs provided by culture (“how things are done here”) allow agents to know what to expect from others’ behavior, thereby facilitating mutual adaption into a new equilibrium. In simple, shared principles/beliefs act as “focal points” that facilitate coordinated adaptation among agents. Li et al (2021) provides supportive empirical evidence exploiting the COVID shock. A limitation of this idea is that change is not only exogenous but is also immediately and effortlessly understood by all agents involved. In the real world, however, change is quite often far from obvious. Organizational members may not know whether the signals they are receiving (e.g., a drop in sales, a change of customer type, a drop in profitability, worse customers reviews) point to a change in the environment or it is simply noise. Answering the question “do we need to adapt?” is not easy as it is endogenous to hard-to-falsify beliefs about the firm’s function. Our model provides a role for culture in answering this question: more accurate beliefs about the function of the organization allow for better signal interpretation and thus, to choose innovation when it is mostly required.[4]

Regarding pro-social norms, our offered contribution is that contrary to received wisdom, but consistent with recent efforts (Hergeux et al, 2021; Chatman et al, 2019), prosocial norms are not monotonically associated with higher performance; instead, they can reduce performance by compromising the capacity of the organization to adapt to changes. A big literature has developed showing that pro-sociality—understood broadly as organizational members helping one another (Barnard, 1937; Fehr, 2018; Rand and Nowak, 2013; Francois et al, 2018; Henrich and Muthukrishna, 2021), favoring the collective instead of the individual (Chatman et al, 2019; Akerlof and Kranton,

[4] There is a third way in which culture can facilitate adaptation without necessarily sacrificing strength (and thus, coordination and efficiency), which is the idea of culture as a toolkit (Swidler, 1986; Howard-Grenville et al., 2020); that is, people have more beliefs and values than those currently being displayed/used and thus can rely on that deeper repertoire to adapt when change strikes. Lix et al. (2022) provide evidence in favor of this view, showing that teams can modulate their “group cognition” to fit the demands of different tasks: more diverse when innovation is required, and less diverse when coordination is required.

2005), and being a good organizational citizen (Organ et al, 2005)—increases member's effort towards organizational goals and thus, increases the organization's performance. However, Chatman (2019) indicates that "examining the empirical research more closely, however, reveals that the seemingly obvious relationship between collectivism and group performance is neither straightforward nor entirely robust […] [and that] despite decades of research, researchers still do not fully understand how collectivism influences group performance" (p. 236-236). Chatman et al (2019) argues, and empirically shows, that the downside of collectivism—which they define equivalently to how we define pro-sociality—is that it can limit performance in certain tasks that require constant adaptation to changing circumstances; they point at pro-sociality blurring perceived individual differences and thus reducing the capacity to draw from diverse ideas, skills, and backgrounds. We complement Chatman et al (2019) in emphasizing the costs of pro-sociality in terms of reduced adaptation. In our model, we fully blur differences—all agents are the same—and thus, we add an additional layer: if sufficiently intense, pro-sociality can be detrimental to performance because the gains from a more efficient tradition are not enough to compensate the losses in adaptation produced by the crowding out of innovation.

2. Literature background

2.1. Formal models of organizational adaptation and organizational culture

Organizational adaptation is a concept with a long and important history (Lawrence and Lorsch, 1958; Cyert and March, 1963; Chakravarty, 1982; Levinthal, 1997; Sarta el al, 2021)[5]. The Carnegie School (CS) tradition has used formal models to study it (see Puranam et al, 2015 for an overview). We share with CS the emphasis on: bounded rationality; the imperfect representations of the underlying reality; learning as behavioral tool to navigate and adapt to an uncertain world; a focus on how the population evolves out of simple rules of behavior; exploration triggered by performance feedback (i.e., problemistic search); and the emphasis that, at the core of organizational adaptation, there is a tension

[5] The review by Sarta et al. (2021) defines organizational adaptation as "intentional decision making undertaken by organizational members, leading to observable actions that aim to reduce the distance between an organization and its economic and institutional environments" (p. 44).

between the "exploitation of old certainties" with the "exploration of new possibilities" (March, 1991; p. 71). However, our model, and cultural evolution models in general, differ from the CS tradition in that we focus on a constantly changing environment instead of a fixed but unknown, and potentially complex, environment; in the CS, the environment can change (and some researchers have explored that), but it is not the central/defining feature of the model. Also, in addition to simulations, which are central to the agent-based models of the CS, cultural evolution models also rely on analytical solutions to a model's equilibrium using the concept of Evolutionary Stable Strategies. The closest model in the CS to ours is March (1991) (and its follow-up, Fang et al, 2010). As we discuss in more detail below in section 4.2, our model differs from March (1991) in that the trade-off between tradition and innovation (or, exploitation and exploration) arises from a different source.

Regarding previous models of organizational culture, an important share of the work has been done in Organizational Economics (Hermalin, 2012 for a review). Some notable efforts are: Cremer (1992), who focuses on the role of common language, knowledge and routines; Van den Steen (2010a), who studies the emergence of shared (and homogeneous) beliefs via self-selection of workers; Gibbons et al. (2021a), who studies how shared cognition in the form of mental frames constrains or enhances the performance level that agents can attain; and Kreps (1990) and Gibbons et al (2021c), who view the repeated-game equilibrium as a shared understanding of "how things are done," that is, general principles/beliefs about proper behavior. We differ from these models in that they assume rationality and use optimization or (classical) game-theory methods (while we assume bounded rationality), and they focus on the homogeneity of beliefs arising from heterogeneous agents while we mute that aspect by default. The models by Dessein and Prat (2022) and Besley and Persson (2022) are closer in spirit to ours: while retaining optimization assumptions, they add dynamics and evolutionary elements (e.g., a birth and death process) to their analysis.

2.2. Two basic dimensions of organizational culture

Schein's influential book, *Organizational Culture and Leadership* (2010) defines culture as "a pattern of shared basic assumptions that was learned by a group as it solved its problems of external adaptation and internal integration, that has worked well enough to be considered valid and, therefore, to be taught to new members as the correct way to perceive, think, and feel in relation to those problems" (p. 17). According to this view, organizational culture has two essential dimensions (namely, *external adaptation* and *internal integration*), which are sufficiently distinct and rich to consider and analyze separately (as Schein does in his 2010 book). All organizations face the following two generic problems, which are reflected in their culture: identifying and understanding their "function for society" (external adaptation – what problems they solve) and "how they come together as a group" (internal integration – social rules). We call the former "beliefs" and the latter "pro-social norms".

2.2.1. Shared beliefs (of varying accuracy)

Beliefs related to external adaptation are those that define a "basic sense of core mission, primary task, or 'reason to be'" (Schein, 2010; p. 89), plus all the beliefs about aspects, as conceptualized by Schein, that derive from having a mission, namely: goals ("here, we believe X is the crucial goal"), means ("here, we believe that in order to achieve X, doing Y is crucial"), measurement ("here, we believe on a specific criteria Z and a specific review process to track progress"), and correction ("here, we take action W when V happens")[6]. The beliefs in this dimension are similar to what Kocak and Puranam (2022) refer to as "causal codes," which capture "beliefs about how the world works" that "causally associate actions to consequences" (p. 8). This way of understanding beliefs of agents is also similar to the idea of the "theory of the business" in some recent work in strategy (Drucker, 1994; Zenger, 2013; Felin and Zenger, 2017).

In contrast to previous formal models of culture—which have tended to focus around the emergence of shared beliefs from heterogeneous individuals—in our model all individuals share the same beliefs (i.e., culture is maximally "strong") and then we vary the verisimilitude/accuracy of these beliefs. To justify varying belief accuracy, we rely on the organizational adaptation literature that has shown that:

[6] Most of the elements in "means" and "measurement," could be labelled as "practices" amenable to formalization and codification (e.g., functional practices in sales, finance, HR). However, these are "culturally" produced in the sense that, as highlighted by Gibbons and Henderson (2013), they do not necessarily rely on formal contracts. Instead, due to many of the subtleties involved in creating them, these elements frequently rely on informal understandings rooted in principles and beliefs.

i) organizations vary in their capacities to adapt (Ahuja and Katila, 2003; Teece, 2007; Helfat and Martin, 2015), ii) an important determinant of this capacity is the cognitive representations (i.e., beliefs, mental models, dominant logics, cognitive frames) that managers, and the organization in general, has about the environment, about itself and what it is capable of (Barr 1998; Danneels, 2011; Eggers and Park, 2018; Eggers and Kaplan, 2009; Helfat and Peteraf, 2015; Barr et al., 1992; Tripsas and Gavetti, 2000; Tripsas, 2009) and that iii) cognitive representations differ in their degree of accuracy or verisimilitude (i.e., the degree to which they reflect correctly the underlying reality), and the more accurate/verisimilar they are, the better for adaptation (Eggers and Park, 2018; Barr et al., 1992; Tripsas and Gavetti, 2000)[7]. The latter is consistent with the Carnegie School view of bounded rationality, in which agents have an imperfect representation of the world that they update via learning over time and thus, can have different degrees of verisimilitude or accuracy (March and Simon, 1958; March, 1991; Puranam, 2018)[8].

2.2.2. Prosocial norms (of varying intensity)

The cultural dimension of internal integration refers to all aspects that determine how the group comes together (Schein, 2010), such as language (i.e., concepts and categories the group uses), boundary and membership (e.g., what defines group membership, necessary/legitimate credentials), power/status (e.g., who is admired), norms regarding cooperation and mutual help (e.g., level of cooperation-prosociality vs. individualism-competition), rewards and punishment (e.g., what informal rules regulate salary levels and changes). Kocak and Puranam (2022) capture this dimension of culture in their idea of “moral codes,” which “assign moral meaning to actions or their consequences … especially considering other regarding preferences” (p. 9).

[7] For example, Eggers and Park (2018) indicate in their review that “cognition and identity *incompatible* with a new technology hinder the acquisition of new knowledge or assets” (p. 363; emphasis added) but also that “*adequate* cognitive frame or vision of management, however, can also reverse the firm’s fate and lead to successful adaptation” (p. 368; emphasis added)

[8] As Puranam (2018) states: “Cognitive representations are conceptual structures in individuals’ minds that encapsulate simplified understanding of the reality these individuals face […] However, representations are *imperfect* [and] […] learning may lead individuals to learn the *wrong lessons* from experience and fall victim of *superstitious learning*. […] However, representations held by an agent need not be perfect (even if that was ever possible) to be useful. They may be *better than*, or at least not worse, than that of others” (pp. 27, 29, 30; emphasis added).

In our model, we focus on one central element of this dimension, namely the norms around cooperation and helping other group members, prioritizing the group rather than the self, and displaying honesty, trust, and goodwill towards group members (Fehr, 2018; Organ et al, 2005; Bolino and Grant, 2016; Rand and Nowak, 2013; Bowles, 2009). A central focus of the literature is trying to explain why some groups display intense prosocial norms and behavior. There are many theories and mechanisms that can support pro-sociality within groups: for example, Francois et al (2018) points at the intensity of competition, Barnard (1938) points at leadership, and Nowak and Sigmund (2005) point at reputation (for reviews, see Rand and Nowak (2013), Bolino and Grant (2016), and Henrich and Muthukrishna (2021)). Thus, it is natural to focus in our model on the intensity of these pro-social norms: we exogenously vary their intensity and focus on the downstream consequences in terms of behaviors, adaptation and performance.

3. The model

We base our model on a workhorse model of cultural evolution labelled "selective learning" (Boyd, 2018; Boyd and Richerson, 2005 chapters 1 and 2). This is a model of adaptation to changes via learning.

Our model revolves around three endogenous variables that capture behavior and adaptation, and two culture parameters that affect our variables. The three endogenous variables are: a knowledge/technology stock, denoted by r; the extent that learning is guided by innovation vs. tradition, regulated by d; and help in diffusing tradition, denoted by a. The two culture parameters are: the "beliefs" about the function the organization plays in the environment, denoted by λ; and "prosocial norms," denoted by u.

When introducing these three variables and two parameters, and the mechanics of the model, we will use a running example to better convey the model's intuition. We will use the example of a salesforce, where individual salespersons sell goods and services to customers, have a selling technique/technology that they obtain via learning, receive a performance signal, and face a changing environment. In table A.1 of the online appendix, we provide an integrated view of the example.

Our model applies to organizations because we assume that, although it may not be easy, parameters λ and u can be manipulated and affected in a top-down fashion. Influencing these parameters is a fundamental aspect of CEOs' and top management teams' role (Schein, 2010; Barnard, 1938; Hermalin, 2012). We think this is feasible in groups up to a certain scale, such as teams, divisions and firms/corporations. In larger units—such as cities or countries—history, politics, and institutions tend to play a much larger role than leadership.

3.1. The environment, learning a technology and shared beliefs

Throughout the model, we assume that the organization is large; to simplify the exposition, we assume it is populated by a continuous of agents with measure N. In our running example, this is a group of salespersons.

There is a large number of possible environmental states. In each period, the state may change with probability p. For every state, there is a unique knowledge set, which we call "technology," that provides higher benefits to agents. Technology captures all the knowledge that agents require in order to be successful in a particular state; it includes skills, tools, routines, practices and processes, all of which may be tacit or explicit. By a normalization, we can assume, without loss of generality, that if an agent carries a technology that is tuned to match the state, then it generates a benefit of 1, and if it is not tuned, it generates a benefit of 0. Generally, these changes in the environment capture many forces that may leave an extant technology at a relative disadvantage, such as disruption by new entrants, a more successful strategy by competitors, changes in the demographic or political environment, changes in the climate, pandemics, and the like.

Organizational adaptation corresponds to the adoption of the tuned technology by agents. At the organization level, the variable r captures the share of agents in the organization that have a technology that is tuned to the state of the world. We call this variable the organization's "stock of adaptive knowledge." In addition to a stock of knowledge, adaptation is also reflected by how well the organization responds to a change in the environment.

In our salesforce example, an environmental state is captured by customers' tastes and preferences; the presence, type and behavior of competing salesforces (e.g., remote/in-person/hybrid); the technology that is used by the actors in this market/industry (e.g., use of digital tools); the presence of substitutes; among others. One or many of these elements may change from one period to another, say between years. The "technology" that salespersons have includes the sales technique they use, the amount/type of preparation they do before approaching a customer or a sales opportunity, the techniques used to scan opportunities in the market, pricing strategy, etc. Depending on the state of the environment on a given year, a particular sales technology is superior, which provides an advantage with respect to all other technologies in terms of performance (e.g., revenue growth, market share, mark-up) (and thus, we can normalize the benefits of technology to 1 and 0 for superior technology and others, respectively).

Agents in the organization obtain technology via learning. In particular, they can resort to two types of learning behavior: individual learning or social learning. Individual learning requires investing time and effort to study and understand the environment which allows the agent to discover the technology that is tuned to the current environment. This learning behavior relies on different tactics such as rational deduction, systematic experimentation, or trial and error. Individual learning has a cost C, which is bounded between 0 and 1 ($0 < C < 1$). Social learning consists of observing what a randomly chosen member of the organization did in the previous period and then copying its technology. Social learning has a cost c, which is bounded between 0 and C ($0 < c < C$). This strategy is less costly because the agent does not need to invest time and effort to understand the underlying environmental state. Thus, these two behaviors display a trade-off for agents: whereas individual learning allows an agent to "get it right" by paying a high cost, social learning is cheaper but bears the risk of a change in nature that renders its technology obsolete. To allow social learning to evolve, we assume that $C - c > p$ (see online appendix A.4.1 for details).

Individual learning can be understood as innovation, without which the organization cannot survive changes in the environment. Social learning can be understood as conformism or tradition, and it is crucial to diffuse adaptive technologies from individual learners (or "innovators") to the rest of the organization. Innovation allows the organization to adapt to changes, and over time, in combination with diffusion via social learning, the organization builds a stock of adaptive knowledge.

With benefits and costs defined, we can start specifying the fitness that the agents within the organization receive. If the state of the world remains unchanged, the fitness an agent obtains with individual learning is $1 - C$ and with social learning is $r - c$. Social learners obtain the right technology, and thus the fitness of 1, only when they copy tuned agents (which occurs at the rate r). If the state of the world changes, the fitness of individual learning is $1 - C$ and of social learning is $-c$. Before we can write down a formal expression for fitness, we first must understand how agents decide their learning behavior.

Agents in the organization choose between individual and social learning after observing a signal s, which is drawn from a distribution that depends on whether the state of the world has changed. This signal can be understood as any feedback the organization receives from its interaction with the environment, such as operational and financial data, customer feedback, or investor reactions. Models of selective learning based on signals are well established in the literature on cultural evolution (Boyd and Richerson, 2005; Boyd, 2018). We follow this literature and model selective learning in a simple way. If the state of the world has not changed, the signal s is distributed exponentially with parameter $\lambda = 1$, which we denote as $f_{\lambda=1}(s)$; if the state of the world has changed, then the distribution is exponential with parameter $\lambda < 1$, which we denote as $f_{\lambda<1}(s)$ (we use $f_{\lambda}(s)$ to denote the density function and $F_{\lambda}(s)$ to denote the cumulative function). Assuming an exponential distribution means that $s > 0$.[9] Given that $\lambda < 1$, when the environment changes, larger values of s are more likely to come from $f_{\lambda<1}(s)$ rather than $f_{\lambda=1}(s)$ (the exponential function flattens out when λ goes down). Thus, large values of s are interpreted by the agent as something "out of the ordinary", which is indicative of a change in the environment, while low values of s are indicative of "business as usual". In our salesforce example, consider a salesperson who, for every period, tracks closely the number of lost sales

[9] Our results hold for any probability distribution with positive support, that fulfills the monotone likelihood ratio property condition; see online appendix A.2. However, our proof of uniqueness of the equilibrium relies on the exponential assumption.

opportunities (say in a B2B context). If a competitor has created a much better product or service (i.e., the world has changed), then the likelihood of observing a large number of lost sales opportunities would increase. If the salesperson observes too many lost opportunities, it is likely that this will be interpreted as a change in the world (i.e., the realization of s is more likely to come from $F_{\lambda<1}(s)$ than from $F_{\lambda=1}(s)$); thus, it may be a good idea to innovate and re-think aspects of the sales technology that he/she uses.

The closer λ gets to zero, when s grows in value, the difference between $f_{\lambda<1}(s)$ and $f_{\lambda=1}(s)$ becomes larger; thus, a lower λ improves the informativeness of the signal.

Notice that the value λ correspond to the agents' capacity to interpret signals that come from the interaction between the organization and the environment. A lower value of λ may come from different sources (e.g., a dedicated market intelligence unit or the education of agents). In this paper, we focus on the beliefs the agent has about the function of the organization plays in the environment—that is, the role and purpose of the organization, as well as all the associated means–ends relations (Schein, 2010; Kocak and Puranam, 2022; March, 1991). It is important to emphasize that these beliefs are *not about changes in the environment*, but about the "structure" of the world—what audiences/customers want, what players believe and how they behave, how profits are made in this setting, what norms are in place, etc.—and what problems the organization is solving and how (the role it is playing in this "structure"). For example, Disney Inc. has a particular understanding of the world, the audience, workers, animators, etc., as well as an understanding of its role, purpose, and associated means–ends relationships (e.g., their business model); this has been documented by Argyres and Zenger (2012), Zenger (2013) and others.

These beliefs can differ in their verisimilitude; that is, they can be a more or less accurate reflection of reality (Puranam, 2018; March and Simon, 1958; March, 1991). An organization can have naïve or accurate beliefs about the "structure" of the world, about its purpose/role in it, and how to fulfill that purpose (i.e., beliefs about the action-consequence pairs, usually embedded in a business model). Drawing from the literature on organizational adaptation (Tripsas, 2008; Tripsas and Gaveti, 2000;

Eggers and Kaplan, 2013; Eggers and Park, 2018) we assume that when these beliefs are more accurate, the interpretation of signals improves, and thus, firms are able to recognize if the environment has changed (e.g., change in customers/audiences' values and preferences).[10]

Continuing with the salesforce example, salesforces have different beliefs about customers: some beliefs are simplistic and inaccurate (e.g., believing that "all customers are driven exclusively by price" in a market with vertical differentiation) which more often than not will lead to signal misinterpretation (such as not realizing that a drop in sales was caused by many customers increasingly valuing quality over price); some beliefs are more accurate (e.g., "some customer prioritize price, but others prioritize quality and service"), improving signal interpretation. Salesperson would also have beliefs about the function of the salesforce: some simplistic and inaccurate (e.g., believing that "we simply take sales orders because customers are captive" in a market with intense competition) which will often lead to signal misinterpretation (such as not realizing that a drop in sales was driven by a new service-oriented strategy by a competitor), or more sophisticated and accurate (e.g., "we provide solutions—a bundle of the right products, the right price, and a great good service—to the customer's unique problem").

The parameter λ is the same for all members of the organization, and thus it is a maximally strong culture: everyone shares the same belief—the same "story" about the organization's function in the world. This allows us to focus in our model on varying the accuracy/ verisimilitude of beliefs. (In the discussion, we return to beliefs heterogeneity within the organization.)

In this setup, how do agents decide their behavior after they receive the signal s? As is standard in Bayesian decision theory, the optimal strategy is a cut-off strategy in which agent i will choose individual (social) learning if the signal s is above (below) a threshold value d_i (Boyd and Richerson, 2005)[11]. This threshold indicates how innovative or traditionalist the agent is: The higher (lower) the

[10] The classic example of Polaroid is enlightening: Polaroid, although it held correct/accurate beliefs—and developed the corresponding knowledge and technology—about the future of imaging (i.e., digital), it had inaccurate beliefs about how profits are made (i.e., profits coming from printing, rather than the camera) and thus the problem and value digital imaging really solves (i.e., physical storage of memories instead of digital) and how to solve it (i.e., maintain marketing and salesforce choices). These inaccurate beliefs impeded Polaroid to correctly interpret a myriad of signals coming from sales, customer reviews, salesforce feedback, etc.

[11] The optimal behavioural strategy is a threshold because the relative advantage of individual learning compared to social learning is monotonic in the signal, therefore there is only one signal level for which both behaviours are equivalent. Above

value, the greater the likelihood that the agent will choose social (individual) learning and thus display traditionalism (innovation). Every agent has a unique threshold. In our salesforce example, salespersons have rules of this type: “If year-to-year sales drop more than XX%, then I will innovate (e.g., try new sales techniques, reroute clients, change discounts and promotions); if not, I will follow tradition (i.e., I will use the “technology/practices/skills” that my peers have been using in the past)” (and XX varies across salespersons).

Now we can write down a fitness expression for agents. Given a level of d_i, the expected individual fitness of an agent if the state of the world does not change is:

$$h_i^u(d_i, r) = F_{\lambda=1}(d_i)(r - c) + (1 - F_{\lambda=1}(d_i))(1 - C) \quad (1)$$

If the state of the world changes, the fitness is:

$$h_i^c(d_i, r) = F_{\lambda<1}(d_i)(-c) + (1 - F_{\lambda<1}(d_i))(1 - C) \quad (2)$$

Therefore, the expected fitness of agent i will be

$$h_i^e(d_i, r) = (1 - p)h_i^u(d_i, r) + p \cdot h_i^c(d_i, r) \quad (3)$$

In our salesforce example, payoff h can be understood as the added value that the salesperson generates for the organization (e.g., price minus cost C or c)[12].

Later in the paper, in the section 3, we explain how the variable d_i evolves and tends toward an equilibrium value for all agents. Notice as well that our model, allows to decouple performance (fitness h) and adaptation (variable r) as different constructs. The organizational adaptation literature so far has tended to conflate these two, leading to some complications (Sarta et al., 2021).

It is important to clarify two elements: information set of agents and types of innovations. First, agents only know about their own d_i, the signal s_i they receive, and their beliefs (which they share with others). They don’t need to know about the probability function F or the value of λ; that is, they don’t need to know about the accuracy of their beliefs. This is intuitive because complex beliefs (as is the

this threshold individual learning is advantageous, and below, social learning is. See online appendix A.2 for a detailed mathematical treatment.

[12] While we don’t model value capture for the firm and employee (i.e., compensation plan of salesperson), we nonetheless can assume that part of this added value will be captured by the salesperson in the form of a compensation and that this share will increase with added value.

case with the “role/function” of an organization firm and its associated “means–outcomes” beliefs) are hard—if not impossible—to falsify. Second, innovation may happen in two circumstances: when the state of nature hasn’t changed, and when it has. In the former case, the agent that receives a very large signal decides to innovate, only to find that the technology it develops is the same as—or different, but not superior to—many others already present in the organization and that they could have obtained via copying. In the latter case, when the world has changed, innovation occurs when it is fully necessary. We label these two types of innovation “ill-timed” and “timely,” respectively.

3.2. Knowledge, technology and practices

The share of agents with tuned technology r (i.e., the stock of adaptive knowledge of the organization) is determined endogenously within the organization as a function of d and the state of the world. To fix ideas, let us assume that all agents in the organization have the same behavior d. In such a case, we can formalize the equation that rules the evolution of the stock of knowledge r. When the state of the world does not change, the evolution of this variable is given by the following difference equation between periods t and $t+1$:

$$r(t+1) = [1 - F_{\lambda=1}(d)] + F_{\lambda=1}(d) \cdot r(t) \tag{4}$$

This equation shows that knowledge stock in $t+1$ depends on the share of agents that innovate, plus the share of conformists multiplied by the knowledge stock in t. When the state of nature changes, only those that decide to do individual learning end up being tuned to the state; therefore, the equation is as follows:

$$r(t+1) = [1 - F_{\lambda<1}(d)] \tag{5}$$

Therefore, the expected level of the stock of adaptive knowledge conditional on the current stock is equal to

$$r^e(t+1) = p[1 - F_{\lambda<1}(d)] + (1-p)\big([1 - F_{\lambda=1}(d)] + F_{\lambda=1}(d) \cdot r(t)\big) \tag{6}$$

3.3. Cooperation and prosocial norms

A large body of literature on organizational learning has documented that the members and productive units of an organization learn from one another, harnessing one another’s experiences to improve practices, processes, routines, and productivity (Argote and Miron-Spektor, 2011; Argote et al, 2021). However, the transmission and sharing of knowledge/information inside the organization is not free of hazards. Knowledge and information are non-rival, non-excludable goods and thus subject to problems of free-riding and under-provision. Therefore, the quality and extent of knowledge/information transmission is not a given, because agents may not exert effort so that others learn better from them (e.g., Sandvik et al., 2020). To tackle this problem, organizations may need to actively promote prosocial norms that support sharing of knowledge/information (Fehr, 2018; Morrison and Wilhem, 2004). Creating conditions for knowledge to be shared is a core function that organizations perform (Kogut and Zander, 1996; Grant, 1996; Brahm and Poblete, 2021), as well as a crucial dimension of corporate culture (Fehr, 2018).

To incorporate these ideas into our model, we add: i) a cooperative effort that agents can exert in order to facilitate social learning by others which we label "help in diffusing tradition", and ii) the prosocial norms the organization holds about help, cooperation and the importance of the group (see section 2.2.2). Help in diffusing tradition is an endogenous variable; the prosocial norms are a parameter. Agents observe prosocial norms.

We first detail the help in diffusing tradition and how it enters the fitness function. An agent can exert effort $a_i \geq 0$ in order to facilitate other organizational members' learning from them—for example, actively sharing insights, mentoring and helping the learner (instead of simply "allowing to be copied"). Agents who are performing either individual or social learning can exert this effort. This helping effort is costly, reducing the agent's fitness h_i in an amount $m(a_i)$; at the same time, however, it reduces the social learning cost of whomever is copying the agent by an amount equal to the agent's effort a_i. We assume that $m(a_i)$ is increasing and convex, and that, for sufficiently low effort levels, $a_i > m(a_i)$, and thus effort a_i is a social dilemma. Formally, we assume that $m'(0) = 0$ and $m'(c) = \infty$ so that it is always socially optimal some interior level of helping effort.

It is important to clarify here that our model only considers innovation brought about by individual experimentation, and not the type of "collaborative" innovation that is produced from the combination of specialized individuals that create something that none of the individuals could do by themselves. Thus, we don't consider how a_i may affect innovation. See section 5.1 for a discussion around this boundary condition of our model.

Consider an organization in which one agent *i* exerts effort a_i, but all the others exert an average helping effort denoted by $\bar{a}$. Now we can expand the expression for agents' fitness, adding helping effort. The expected individual fitness of agent *i* is as follows. If the state of the world does not change, fitness is

$$h_i^u(d_i, \bar{a}, a_i, r) = F_{\lambda=1}(d_i) \cdot (r - c + \bar{a}) + \left(1 - F_{\lambda=1}(d_i)\right) \cdot (1 - C) - m(a_i) \tag{7}$$

If the state of the world changes, fitness is

$$h_i^c(d_i, \bar{a}, a_i, r) = F_{\lambda<1}(d_i) \cdot (-c + \bar{a}) + \left(1 - F_{\lambda<1}(d_i)\right) \cdot (1 - C) - m(a_i) \tag{8}$$

Thus, the expected fitness of the agent will be,

$$h_i^e(d_i, \bar{a}, a_i, r) = (1 - p) \cdot h_i^u(d_i, \bar{a}, a_i, r) + p \cdot h_i^c(d_i, \bar{a}, a_i, r) \tag{9}$$

The expected individual fitness is, therefore, increasing in other members' effort and decreasing in the agent's own effort.

It is useful to measure agent *i*'s contribution to the organization via help in diffusing tradition as the difference between i) the actual fitness of all other agents in the organization when they exert helping effort $\bar{a}$ and ii) the fitness that would exist if the focal agent were to exert no helping effort at all, instead of effort a_i. This contribution by the agent to the organization corresponds to the effort a_i times the likelihood of social learning by other members in the organization. To see why this is the case, consider an organization with *N* members in which everybody has learning rule d (the same for all) and exerts a helping effort $\bar{a}$. The expected fitness of the whole organization is given by $N \cdot h(d, \bar{a}, r)$ (which is equivalent to the "profit" of the organization). If agent i exerts no effort instead of exerting a_i, everybody else's fitness will be decreased by an amount a_i whenever they learn socially from agent *i*. Agents learn socially with probability $(1 - p)\, F_{\lambda=1}(d) + p\, F_{\lambda<1}(d)$ and they copy from agent *i* with probability $1/N$. Therefore, the fitness of the organization becomes

$$N \cdot \{\, h(d, \bar{a}, r) - \frac{1}{N}[(1-p)F_{\lambda=1}(d) + p\, F_{\lambda<1}(d)]\, a_i\} \tag{10}$$

and, thus, the expected contribution of the agent to the organization is

$$\psi_i(d, a_i) = [(1-p)F_{\lambda=1}(d) + p\, F_{\lambda<1}(d)]\, a_i \tag{11}$$

Now we comeback to "prosocial norms" and address their connection to "help in diffusing tradition". We model this by making the contribution of agents $\psi_i(\vec{d}, a_i)$ important and valued in the organization. We expand the fitness h_i^e of agents by adding this contribution, and we label the resulting quantity as "payoff." We represent the expected payoff for an agent as follows:

$$q_i^e(\vec{d}, \bar{a}, a_i, r) = h_i^e(d_i, \bar{a}, a_i, r) + \mu \cdot \psi_i(a_i) \tag{12}$$

A positive μ captures the presence of "prosocial norms" in the organization's culture; a larger μ indicates that these are more intense. The presence and intensity of these norms dictates how much the agent benefits from its contribution ψ_i; that is, if the firm's culture values pro-sociality, then the agent benefits from helping in the diffusion of tradition. Similar to λ, we assume that μ is shared by all in the organization, so that we can focus on the impact of varying its intensity. The survey-based literature on organizational culture has documented considerable variation in how intensively organizations as a whole prioritize the group versus individual (Denison and Mishra, 1995; Chatman and O'Reilly, 2016; Cameron and Quinn, 2011; Graham et al., 2022; Chatman et al, 2019). Organizations can affect the intensity of prosocial norms and norms via several levers, including: symbolic rewards/punishment (e.g., creating status/prestige around helping effort); tying awards to sharing; stimulating peer pressure/punishment to free-riders that do not exert helping effort; formal levers such as establishing and enforcing processes of knowledge sharing (e.g., mentoring schemes); creating pecuniary incentives (e.g., bonuses tied to group performance); and adding formal criteria for performance assessment and promotions.

Returning to our running example of a salesforce, when salespersons learn from one another about how things are done (e.g., sales techniques, daily route planning, priorities in terms products/customers, etc.) they may receive help (or not) from the salesperson that is being sought out as the source of knowledge/tradition. The salespersons that help transmitting tradition to others generate a benefit to others and the organization as whole, but a cost to themselves. Motivation to help is generated if the culture of the salesforce emphasizes pro-sociality in its norms.

Notice that whereas the organization cares about the expected fitness of the whole organization—$N \cdot h_i^e(d_i, \bar{a}, a_i, r)$, which we label "profits" or "performance"—the agents care about expected payoff $q_i^e(\vec{d}, \bar{a}, a_i, r)$. This discrepancy can introduce a tension between the organization and its agents because $h_i^e(d_i, \bar{a}, a_i, r)$ and $q_i^e(\vec{d}, \bar{a}, a_i, r)$ may respond differently to changes in the environment.

4. Solution and findings

In this section, we first present the criteria (i.e., evolutionary stable strategies) and the equations that describe the equilibrium of the model, and then we present the findings we obtain exploring how the equilibrium changes when the cultural parameters λ and u are modified.

4.1. Evolutionary stable strategies

We assume that in each period, a small share of agents copy the traits help in diffusing tradition a_i and learning threshold d_i of the agents who have a higher payoff q_i^e (this is independent of the copy involved in social learning). In evolutionary modelling, this updating process is known as "replicator dynamics": in biology, replication means that fitter organisms have more offspring, in social science, more successful traits or strategies are copied with higher likelihood. While in biology replicator dynamics means that different species grow or decay in number, in social science it means that the behavioural repertoire (in our case d_i and a_i) diffuses or vanishes. As a result of this updating process, over time, agents "converge" on particular set of values for a and d. To determine whether such values exist—and whether they are unique—we rely on the concept of evolutionarily stable strategies (ESS) (Maynard Smith, 1982).[13]

[13] This convergence to an ESS does not mean that change cease to occur. Rather, it means that in the long run, the variables continue to change—albeit doing so around the neighborhood of the set of "equilibrium" or "long run" values for a and d defined by the ESS.

We follow Boyd and Richerson (2005) in defining a symmetric ESS as a situation in which a small set of agents that deviate from a^* and d^* perform strictly worse in expectation than other agents in the organization. The pair a^* and d^* are said to be "stable" because, in the long run, agents copy from more successful behaviors; as a result, deviations will not spread in the organization. (As indicated above, it is not the case that agents have fewer offspring if they deviate from ESS; rather, their behavioral repertoire diffuses or vanishes via replicator dynamics because they have a lower payoff.)

To find the ESS, first consider possible deviations in the effort exerted in helping diffuse traditions. The expected payoff of an agent in ESS is given by $q_i^e(a^*, d^*, r) = h_i^e(a^*, d^*, r) + \mu \cdot \psi_i(a^*, d^*)$. An agent who deviates to an effort level $a^* + \varepsilon$ will achieve a lower fitness h_i^e by exerting more effort helping others, but their contribution to the organization ψ_i increases; these two effects are reflected in its payoff:

$$q_i^e(a^* + \varepsilon, d^*, r) = [h_i^e(a^*, d^*, r) + m(a^*) - m(a^* + \varepsilon)] + \mu \cdot \{\psi_i(a^*, d^*) + \varepsilon[(1-p)F_{\lambda=1}(d^*) + p\, F_{\lambda<1}(d^*)]\} \qquad (13)$$

In an ESS, the expected payoff of the mutation should be less than the payoff of the members of the organization and, therefore, the condition for ESS is $q_i^e(a^* + \varepsilon, d^*, r) \leq q_i^e(a^*, d^*, r)$ for every possible deviation ε. A necessary condition for this is that $\partial q_i^e(a^* + \varepsilon, d^*, r)/\partial\varepsilon$ evaluated at $\varepsilon = 0$ must be 0 because we are in maximum. This is equivalent to the requirement that $\partial q_i^e(a^*, d^*, r)/\partial a_i$ be equal to zero at $a_i = a^*$, which, after reordering, yields the following:

$$m'(a^*) = u[(1-p)F_{\lambda=1}(d^*) + p\, F_{\lambda<1}(d^*)] \qquad (14)$$

This is a standard first-order condition that equalizes the marginal cost of helping effort $m'(a^*)$ to its marginal benefit, which is the weight u times the probability of being copied $[(1-p)F_{\lambda=1}(d^*) + p\, F_{\lambda<1}(d^*)]$.

Analogously, mutations with respect to the learning parameter d^* should not increase the payoff of agents in the organization. Observe that at the maximum, these mutations affect only the fitness h_i^e and not the contribution ψ_i in the payoff of the agent; therefore, d^* satisfies the first-order condition $\partial q_i^e(a^*, d^*, r)/\partial d_i = \partial h_i^e(a^*, d^*, r)/\partial d_i = 0$ at $d_i = d^*$, which yields the following:

$$(1-p)f_{\lambda=1}(d^*)(r-c+a^*-[1-C])+p\,f_{\lambda<1}(d^*)(-c+a^*-[1-C])=0 \tag{15}$$

In equation (15), the first term represents the change in fitness from increasing social learning (and decreasing individual learning in a corresponding quantity) when it doesn't cause harm—that is, when the world hasn't changed; this must be equal to the case in which it does cause harm (the second term)—that is, when the world has changed.

Finally, to calculate the long-term unconditional expected level of the stock of adaptive knowledge, we use the difference equation (6), which shows how this adaptive knowledge changes from one period to the next. The unconditional expected level of knowledge is reached when it does not change from one period to the next—that is, when $r(t)=r^e(t+1)=r$. Using this steady-state condition, we can obtain the long-term level of the stock of adaptive knowledge r:

$$r=\frac{p\left(1-F_{\lambda<1}(d^*)\right)+(1-p)\left(1-F_{\lambda=1}(d^*)\right)}{1-(1-p)F_{\lambda=1}(d^*)} \tag{16}$$

Equations (14), (15), and (16) characterize the ESS. In the online appendix A.3, we provide a proof that the equilibrium exists and that it is unique.

4.2. Trade-off between innovation and help in diffusing tradition

The first result we would like to highlight is the negative relationship between innovation (measured as the share of agents who perform individual learning) and help in diffusing tradition (measured by a). This negative relationship comes directly from equation (14), which clearly describes a trade-off between help in diffusing tradition and innovation. Intuitively, more innovation implies less social learning, which reduces the benefits—and thus, frequency—of helping effort a. Mathematically, if the share of social learning $[(1-p)F_{\lambda=1}(d^*)+p\,F_{\lambda<1}(d^*)]$ increases (for example, due to a more stable environment[14]), then, for a given u, the helping effort a^* would also need to increase in order for the marginal costs to equalize the marginal benefit in equation (14). If the strength of pro-social norms u increases, the situation is the same: social learning would be favored and increase and, thus, a^* would

[14] In online appendix A.4, we detail the impact of changing environmental uncertainty p on the model equilibrium. In online appendix A.5, we also explore how changes in C and c affect the equilibrium.

increase because it has a larger marginal benefit. Figure 1 illustrates this trade-off by graphing how the ESS equilibrium shifts by changing p.

This trade-off suggests that we should observe a tendency for the culture of organizations to focus on either innovation or tradition. This prediction is consistent with the view that emerges from three prominent organizational culture frameworks (Denison and Mishra, 1990; Groysberg et al., 2018; Cameron and Quinn, 2011), which indicate that there will be a tension between innovation and tradition—and in the case of Cameron and Quinn (2011), a trade-off as well[15].

Furthermore, this result is consistent with the literature highlighting the tension that organizations exhibit between exploration and exploitation (O'Reilly and Tushman, 2013; March, 1991; Uotila et al., 2008; Lavie et al., 2010; Benner and Tushman, 2006). However, unlike the model by March (1991), where the trade-off arises due to the presence or lack of diversity of beliefs across organizational members, in our model we assume homogeneity of beliefs, and thus the trade-off arises from a different source: agents have to make a choice in every period, under uncertainty, whether to experiment/innovate (individual learning) or follow tradition (social learning). The trade-off arises in our model from the difficulty of choosing innovation or tradition when change is uncertain—not by the homogeneity v/s diversity of agents as in previous work (e.g., March, 1991; Van den Steen (2010a and 2010b).

[Insert Figure 1 around here]

4.3. Accuracy of beliefs

We split the discussion of the findings in two. First, we discuss the impact on adaptation and performance, and the empirical evidence consistent with our findings, and then we discuss in detail where these results come from, that is, what are the mechanisms of the impact of λ.

[15] Cameron and Quinn (1983) indicate the following: "From the external view, […] the emphasis is on the overall competitiveness of the organization in sometimes changing environments. From the internal view, the organization is a socio-technical system. Participants have unique feelings, likes and dislikes, and require consideration […]. *When the external value on the overall organization is maximized, the internal emphasis on the socio-technical equilibrium may be reduced; and when the emphasis on internal harmony grows, it may tend to shift emphasis away from overall competitiveness*." (p. 370; emphasis added)

To discuss performance and adaptation we rely on Figures 2 and 5. In Figure 2, we illustrate how the ESS of the model changes when "beliefs" in the organization improve from λ = 0.7 to λ = 0.1 (remaining parameters are set as in figure 1 and p = 0.1). By "improve," we mean that the beliefs more accurately depict the structure of the environment (not whether it has changed or not) and the function of the organization in it; this means that the signals produced by the interaction of the organization and the environment are more informative to agents. In the graphs of Figure 5, we illustrate the dynamics of adaptive knowledge (left-hand graph) and fitness (right-hand graph) for different values of λ.

Regarding performance, we find that there is a monotonic increase in fitness as the beliefs become more accurate (see bottom-left graph of Figure 2). This means that the organization will always improve its performance—organizational performance is equal to N times fitness—when its beliefs increase in accuracy. Regarding adaptation to shocks, Figure 5 shows that when λ is lower, adaptation to shocks improves. The left graph of Figure 5 shows that the drop in the stock of adaptive knowledge when the world changes (in period 10) is much smaller, and the same can be appreciated for fitness (right graph of Figure 5).

In sum, more accurate beliefs improve the stock of adaptive knowledge and the response to shocks, and generates higher performance for the organization.

Research on cognition and adaptation cited above in section 2.2.1 (e.g., Barr et al., 1992; Tripsas and Gavetti, 2000; Tripsas, 2009) provides supportive evidence for this result. However, this research takes the form of qualitative, in-depth case studies (e.g., Polaroid in Tripsas and Gaveti, 2000) and we are not aware of any large-scale quantitative studies that have tested this finding thoroughly (most likely because it is notoriously difficult to measure the accuracy of beliefs for many organizations at once). There is, nonetheless, some evidence we could point at, depending on how far one stretches the interpretation of the concepts and measures. First, Costanza et al. (2016), using a sample of 95 firms, found increased odds of survival over the next 70 years for those firms that in 1940 held three values of adaptability, namely: i) "the organization pays attention to its external environments and values reading and interpreting signals from their environments," (ii) "the organization proactively works to identify internal and external problems; the organization focuses on future problems or environment changes," and iii) "the organization believes that it has the ability to change"), provided that these values were backed by the capacity to change. Second, Vicinanza et al. (2021) show that firms capable of better vision (i.e., firms with more accurate beliefs about what will happen) enjoy higher stock returns.

The finding so far is intuitive, and admittedly close to the model's assumptions. Nonetheless, we think there is value in formalizing the intuition that comes from a literature that has been, to date, mostly conceptual in its theory development. We now turn to a more subtle finding, uncovering the mechanisms through which more accurate beliefs generate these adaptation and performance benefits.

This increase in fitness from lower λ comes from three sources: i) an increase in the stock of adaptive knowledge r (top-left graph of figure 2), and iii) an increase in the less costly learning strategy, social learning (top-tight graph of figure 2) without sacrificing adaptiveness, and iii) an increase in helping effort a (bottom-right graph of Figure 2). Perusing figure 2 and equations 7 to 9, these three sources are roughly equal in magnitude (with source iii weighing slightly less than the other two). We unpack each one in turn.

The stock of adaptive knowledge increases because, even though total innovation is decreasing (top-right graph of figure 2), agents time their innovation better because they are more accurate in inferring from the signal whether the world has changed or not. Recall that innovation can happen in two circumstances: when the world has not changed ("ill-timed" innovation) and when the world has changed ("timely" innovation). In the former, the amount of innovation is $1 - F_{\lambda<1}(d^*)$; in the latter, it is $1 - F_{\lambda=1}(d^*)$. The left-hand graph of Figure 3 depicts these two proportions against the value of λ. The figure shows that innovation becomes more efficient—that is, it occurs increasingly when the world changes, and not when it doesn't change. This better timing of innovation is reflected in figure 5, which shows that the drop in adaptive knowledge and fitness when change strikes is smaller when beliefs are more accurate. Innovation also becomes "more productive": in Figure 2, the stock of adaptive knowledge increases (top-left graph) at the same time that innovation is going down (top-right graph), indicating an increase in the amount of adaptive knowledge per innovation unit. In short, adaptive knowledge increases because innovation becomes better-timed and more productive.

In turn, focusing innovation on periods of change allows to increase the use of social learning when it is most useful, that is, during periods of stability. As "ill-timed" innovation goes down (left-hand graph of Figure 3), social learning in times of stability goes up. This leads to an overall increase in social learning (top-right graph of figure 2), a cheaper learning strategy that, given that innovation is well focused, doesn't come with higher organizational liability to environmental changes (which would otherwise happen with a high value of λ). The benefit of more social learning during stability is reflected in the period right after a shock. In figure 5, fitness in period 11 bounces back quicker from period 10 (when the shock occurs) when λ is lower because more social learning transmits the higher stock of knowledge in period 10 quicker within the organization.

Then, in response to more social learning, more help in diffusing tradition is produced (bottom-right graph of figure 2). This increases fitness by reducing the cost of tradition in the organization. As it can be appreciated in equation (14) and the discussion in subsection 4.2, this occurs because the marginal return to help in diffusing tradition increases.

At the heart of these mechanisms, particularly i) and ii), is the fact that more accurate beliefs reduce, or minimize, the trade-off between innovation and help in diffusing tradition. In our model this trade-off arises from the difficulty in choosing a learning strategy—tradition or innovation—under uncertainty. In Figure 4, we plot the relationship between total innovation (left-hand graph), timely innovation (middle graph) and ill-timed innovation (right-hand graph) on the one hand and help in diffusing tradition on the other, as the organization moves from high to low λ (using the same range as in Figures 2 and 3). The middle graph shows that a positive relationship is formed between innovation and help diffusing tradition during periods of environmental change. This occurs because the reduction of λ increases help in diffusing tradition (bottom-right graph of Figure 2) and timely innovation increases (Figure 4). Notice that the trade-off persists for total innovation (equation 14 means that there will always be a trade-off, however small), although it is weaker in comparison to Figure 1; only for "timely" innovation the trade-off is broken.

The image that emerges from this discussion is that an organization with accurate beliefs will be very responsive during a period of change (or disruption), innovating more than an organization of less accurate beliefs during that period—and that, after the new technology is generated, the organization with accurate beliefs is capable of diffusing it more rapidly during periods of stability thanks to more and cheaper tradition. Overall innovation cedes terrain to tradition, but both become more efficient, minimizing the trade-off between them.

Overall, we summarize the findings on the mechanisms of λ with the following proposition.

> <u>Proposition 1</u>: *Increasing the accuracy of beliefs focusses innovation on periods of environmental change, makes innovation more productive, reduces (increases) the frequency of innovation (tradition), increases help in diffusing tradition, and reduces the trade-off between innovation and help in diffusing tradition.*

These mechanisms are not entirely obvious ex-ante, not without a formal model. Also, this proposition provides precise testable predictions for empiricists.

This proposition is consistent with research on ambidexterity, that indicates that organizational culture is one way in which organizations can minimize or break through the trade-off between exploration and exploitation (O'Reilly and Tushman, 2013). In particular, our model is in line with the idea of "contextual ambidexterity" introduced by Gibson and Birkinshaw (2004) in which the organization provides the right context—including culture—that "encourages individuals to make their own judgments as to how to best divide their time between the conflicting demands for alignment and adaptability" (ibid; p. 211). O'Reilly and Tushman (2013) indicate that contextual ambidexterity is "a function of a *culture* that supports workers to pursue exploration and exploitation" (p. 329; emphasis added). However, while this idea emphasizes engaging with exploration and exploitation simultaneously, our model and finding emphasizes the timing of when to explore and when to exploit. This connects with the idea of "sequential ambidexterity" where exploration and exploitation are pursued in different moments in time (O'Reilly and Tushman, 2013), and is in line with the effort by

Randle and Pisano (2021) which indicate, consistent with our model, that the within-firm evolutionary process that generates innovations “begins with a period of exploration, which is followed by a period of exploitation” (Ibid, p. 291).

[Insert Figures 2, 3, 4 and 5 around here]

4.4. Intensity of prosocial norms

As in section 4.3, we split this discussion into two parts. In the first we describe how the intensity of prosocial norms affect helping effort, the extent of tradition versus innovation and adaptation. In the second we describe the impact on organizational performance (fitness). The first is closer to the assumptions we make in the model, while the latter is more unexpected and novel.

In Figure 6, we display how the ESS of the model changes when the intensity of prosocial norms is increased from 0 to 0.6. We find that increasing u leads to a large increase in help in diffusing tradition (see the bottom-right graph of Figure 6). This is to be expected as u increases agents’ payoffs if they help others following tradition better (see equation 14). This promotion of helping effort by the cultural norms generates then consequences for innovation and adaptation. Increasing u provides an advantage to social learning and tradition (it becomes cheaper), and thus, we find in the top-right graph that tradition substitutes for innovation as u grows. The right-hand graph of Figure 3 shows that the reduction in innovation as u increases occurs for both timely and ill-timed innovation. This leads to a reduction in the stock of adaptive knowledge (top-left graph of Figure 6) as well as a reduction in the capacity to adapt to environmental changes when these occur—both graphs in Figure 7 show that the drop in r and fitness when the environment changes is more precipitous when u is larger. All of this is reflected in that changes in u generate a larger trade-off between innovation and help in diffusing tradition (see figure 8) than the trade-off produced when p changes (see figure 2) or when λ changes (see figure 4).

Several lines of research provide consistent evidence with the result depicted so far, namely, when tradition is favored (in our case via prosocial norms that boost help in diffusing tradition), innovation and adaptation tend to suffer (see the review by Sarta et al., 2021). This result is intuitive.

The second result on how u affects performance is non-obvious. As it can be appreciated in in the bottom-left graph of figure 6, we find that fitness—and, by extension, organizational performance—has an inverted-U relationship with u. The highest fitness is attained around $u \approx 0.24$. The intuition for this inverted-U result is that when u is growing, at first it provides benefits to the organization because social learning is not only increasing but more importantly it is becoming cheaper because helping effort grows (bottom-right graph of figure 6) and this reduces the overall cost of tradition in the organization faster than the drop of fitness that comes from having a less adaptive technology due to lower innovation (top-left graph of figure 6); however, when u gets too big, there comes a point when the gains in terms of cheaper, more efficient tradition are eventually outweighed by the reduction on innovation and the stock of adaptive knowledge, and fitness starts to go down. This inverted-U finding is counterintuitive—and even surprising—because the literature, for the most part, tends to posit that stronger pro-social norms in a group are always better (Fehr, 2018; Boyd and Richerson, 2005). Instead, our results suggest that pro-social norms have an important downside that prior literature has largely not considered, namely that of reducing innovation and adaptation. Furthermore, it implies that an ideal level of u exists for organizations. We summarize this finding in the following proposition.

<u>Proposition 2:</u> *For low levels of prosocial norms there is a positive relationship between prosocial norms and performance, but this relationship turns negative after a threshold, creating an inverted-U relationship.*

Regarding empirical evidence, to the best of our knowledge we could find only two papers that document how pro-sociality and cooperation can be detrimental to organizational performance, both showing that it does via may hindering adaptation to changes/shocks. The closest to our paper is Hergeux et al (2021) who studies reciprocity-based cooperation (i.e., cooperating with others if they respond in kind) in the context of open-software development teams. They find that reciprocity in the team has a non-linear with its performance: more reciprocity generates better performance, but also makes the teams more likely to fail. The reason is that while reciprocity can build up and sustain cooperation in good times—team members feeding of each other's effort—, if "something" happens that makes a team member cooperate significantly less (e.g., a personal matter; better projects to work on), then cooperation may unravel because it is conditional. The paper, however, doesn't specify precisely what this "something" is; we contend that it can certainly be environment shocks/changes as we model them in our paper (there could be other reasons though, such as a software developer giving attention to an alternative project that is unobserved by the rest of the team).

The second is Chatman et al (2019) which studies the role of collectivism—defined as "a norm in which the demands and interests of groups are prioritized above individual needs and desires to achieve collective goals" (p. 235) and whose primary/defining feature is "cooperation with relevant group members" (p. 235)—in the context of groups climbing Himalayan mountains. They find that collectivism blurs individual differences, and this hurts the group in tasks in which performance is a function of the most expert member (safety) but helps the group in tasks that all group members need to do (reaching the summit). This study is consistent with our theory: it finds that, in a context of adaptation to changing circumstances, pro-sociality can be detrimental.

Proposition 2 has three corollaries, all amenable to empirical study. The first corollary is that "*there is degree of intensity of prosocial norms that maximizes the performance of the organization*". In a context of competition among organizations—what is called "group selection" in cultural evolution theory (Boyd and Richerson, 2005)—it is possible that over time organizations that have the "right" amount of pro-sociality are going to be selected. Also, this could explain why large firms' persistent struggles to attain large scale cooperation; for example, a survey of 1,348 CEOs of large US firms ranked cooperation among employees as the main driver of an effective culture and 84% believe the cooperation level should be higher (Graham et al., 2022). In contrast to the usual explanations around the difficulty of sustaining cooperation in large groups due to increasing free-riding temptation, it could be that, in many cases, a scaled back pro-sociality is beneficial and thus, selected (independent of the opinions of CEOs).

The second corollary is an extension of the first: "*the optimal level for the intensity of prosocial norms increases with the degree of environmental uncertainty p*". In the online appendix A.4 we provide the analysis that supports this statement. This is consistent with evidence indicating that within group pro-sociality increases when the group faces external threats (e.g., war, natural disasters, pandemic) (Bauer et al., 2016; Gelfand et al., 2011; Francois et al., 2018; Bohm and Rockenbach, 2012; Bowles, 2009). Again, group selection can select organizations that display the "right" amount of pro-sociality for the environmental instability they face.

These first two corollaries suggest that a cultural evolution approach to organizational culture—via group selection—can bring back attention to the environment as a driver of organizational culture. For example, the classic work by Deal and Kennedy (1982) suggests cultures arise from the nature of workflow and the risk/reward in the environment. Consistently, Sorensen (2002) showed that strong culture produces more reliable performance only in stable industries. More recently, Ulrich et al (2009) recommends to build "outside-in" cultures, where customers—and the function the firm/organization performs for them—is central in informing the desired culture. Francois et al (2018) shows that cooperation is higher in firms that operate in more competitive settings. In a way, by emphasizing culture as the repository of beliefs that proved successful/adaptive to external challenges, Schein (2010) is also part of this camp.

The third corollary of proposition 2 is "*more intense prosocial norms always increases the payoff of organizational members and thus, creates a tension with the organization after the optimal intensity level*". This can be appreciated in the bottom-left graph of Figure 6. This difference is produced because agents receive benefits in the form of recognition and rewards beyond fitness (see equations 11 and 12). This misalignment between agents and the organization suggests that there might be risk of prosocial norms becoming too strong if they are left to emerge spontaneously from agents' interactions (and group-selection is weak). We are not aware of studies documenting this phenomenon.

[Insert Figures 6, 7 and 8 around here]

4.5. Complementarity between causal and prosocial norms

In Figure 9 we display the impact on fitness by varying λ and u simultaneously. We find that the accuracy of beliefs and the intensity of prosocial norms are complement; that is, if one increases, the impact on fitness and organizational performance improves more (or is less detrimental) if the other increases as well. We can see this result when comparing the slopes of the two graphs of Figure 9. The mechanism that drives this result is that making innovation more efficient, in the sense of innovating when it is needed (driven by more accurate beliefs) will be even more useful when knowledge is spread internally more effectively, so that the new adaptive technology spreads faster after an environmental change has occurred. The same occurs in the other direction; the value of spreading knowledge internally increases when the knowledge being spread is better tuned to the state of nature.

Overall, we summarize this discussion in the following proposition.

<u>*Proposition 3*</u>*: The intensity of prosocial norms and the accuracy of beliefs are complementary.*

[Insert Figure 9 around here]

As with proposition 1, given the difficulty in measuring the accuracy of beliefs, quantitative evidence is hard to obtain. An alternative is to look at studies that measure the mechanism: knowledge sharing is more impactful if the innovation being shared is better, and vice-versa, improving innovation is more impactful if the sharing of the innovation is more efficient. The setting of franchise systems is a good setting for this: Watson et al (2014) uses a large sample of UK franchisees (across many sectors) and shows that i) at least 40% of franchisees innovate locally—a new product, service or process created in a specific franchisee—and ii) there is heterogeneity whether these innovations are shared across franchisees that share a franchisor. Then, they do in-depth interviews on 29 franchisees and show that cooperation and good relations among franchisees increase local innovation and the sharing of these innovations among franchisees. Several quotes suggest that this is due to sharing and innovation being complements: both are better together. The papers by Colla et al (2019) and Dada et al (2012), which use quantitative and qualitative techniques respectively, are broadly consistent with the findings of Watson et al (2014).[16]

[16] Additional indirect evidence for the mechanism underlying proposition 4—albeit less precise than the franchise research—is the ambidexterity literature which shows convincingly that firms that are able to exploit and explore perform better (O'Reilly and Tushman, 2013).

5. Discussion

5.1. Extensions to the model

All models make simplifying assumptions about the world. We would like to highlight four assumptions in our model and speculate how lifting them may yield novel findings. We believe these are fruitful avenues for future research.

The first is that, because technology travels between individuals in our model, it applies well to: i) organizations where individuals are the productive units performing a similar activity, such as service industries (law, consulting, accounting, medicine, education, etc.); ii) sub-units within organizations that rely on individuals executing a similar activity (e.g., sales organizations, customer service, logistics); or iii) organizations comprising a collection of productive units, such as franchises, stores or plants, which can be interpreted as "unitary entities" on their own, and can thus be treated as the agents in our model for analytical purposes. However, this excludes many organizations whose output is an outcome of a web of collaborative effort across many individuals and units; consider, for example, the assembly of complex capital goods such as airplanes. This would require quite a different modelling approach, which is still in its infancy in the field of cultural evolution (Smaldino, 2014). In these complex product/services scenarios, innovation is a process that involves many individuals and units coming together to produce something they would not be able to produce individually. One simple way to start exploring this is to allow for the return to innovation to be increasing on the share of innovators; this might capture productive collaborations between specialized innovators.

A second assumption is that the internal structure of the organization is absent in our model; that is, all agents are undifferentiated and can connect with one another with the same probability. A good way to lift this assumption is to allow social learning to be non-random, so that agents learn disproportionately from a subset, such as the most prestigious, or the majority. The literature on cultural evolution is deep in the modelling these "biases" in social learning (Boyd and Richerson, 2005; Kendal et al, 2018).

A third assumption is that culture is homogeneous; that is, all agents have the same λ and u. This allows us to "shut down" the issue of cultural strength and instead focus on how culture affects adaptation. However, it would be very interesting to allow these dimensions to vary across individuals. One could explore more nuanced questions regarding the relationship between cultural strength and organizational adaptation. For example, research on "groupthink" warns that homogeneity of beliefs can lead to worse outcomes (Janis, 2008) and on "group polarization" shows that groups in general generate more extreme beliefs/opinions than its individuals (Isenberg, 1986; Friedkin, 1999; Zhu, 2013). Therefore, one could add a cost of homogeneity of beliefs alongside benefits of heterogeneity (March, 1991).

A fourth assumption is that of the "stability of beliefs." In our model, λ does not change over time. We argue that when λ is smaller, environmental changes are easier to identify and innovation is triggered when most needed, improving the stock of adaptive technology r; given that tradition and help diffusing it are promoted by a smaller λ, the technology that was discovered by well timed-innovation diffuses quicker inside the organization. However, given a change in environment, and thus a change in technology r, it is not obvious that these cultural beliefs can be stable. In order to be stable, our model assumes either that i) λ captures beliefs at a high level of abstraction such that they do not need to change (at least in short/medium term) in order to be "applicable" to any environment—similar to the view by Gibbons et al (2021c)—, or that ii) λ can actually change over time in a seamless, costless way. Both assumptions have different pros and cons. To lift these assumptions, one could model the adjustment process of beliefs in response to environmental changes (perhaps following the ideas from the Carnegie School; see Puranam et al, 2015); also, one could create a hierarchical structure of beliefs such that a subset of beliefs "at the top" do not change, whereas a majority of lower-level and "closer-to-the-action" beliefs do need to change to accommodate to a new environment (this subset can actually be a part of r, as many "practices" are actually cultural, not formal elements; Gibbons and Henderson, 2013; Howard-Grenville et al., 2020). Empirical inspiration for these efforts can be taken from the literature on narratives and stories, which has pointed to their importance for organizational change both in providing stability and in enabling change (e.g., Vaara et al., 2016; Dal Piaz and Di Stefano, 2018; Gibbons, 2021a).

5.2. Broader implications

Two broader implications emerge from our model. First, our model connects the organizational culture and organizational adaptation literatures. These two have been disjointed so far: if one makes a word-search of “culture” in the reviews of adaptations by Sarta et al. (2021) and Eggers and Park (2018) yields zero hits. This is unexpected as cognition plays an important role in scholarship on both adaptation (Barr et al., 1992; Tripsas and Gavetti, 2000) and organizational culture (DiMaggio, 1997; Hannan et al., 2019; Schein, 2010). We believe that further connecting these two literatures, in particular using insights from organizational adaptation, for example the idea of “cognition/belief accuracy” (in our case) and “broad stories/principles as guides” (in the case of Gibbons, 2021c), can be very fruitful to understand how culture can facilitate adaptation. The second implication is that our model may be used to formalize problemistic search (Cyert and March, 1963). Empirical research on problemistic search has expanded, but the “development of the theory has not kept pace” (Posen et al, 2018; p. 208). We believe cultural evolution models that rely on signals, such as the selective learning we use in this paper, can be used to formally model problemistic search. Consider that in problemistic search “a firm’s recognition of performance below aspirations, which is the level of future performance deemed acceptable, leads to a process of search to discover a solution to the problem” (ibid). Our model is consistent with this description, and we believe future research could explore this connection.

6. Conclusion

In this paper, we answer calls to study the mechanisms that drive the impact of organizational culture on organizational performance (O’Reilly and Chatman, 2016; Chatman and Srivastava, 2021; Chatman et al., 2014; Gibbons et al., 2021b; Grennan and Li, 2022). In this paper, we study how organizational culture can affect organizational adaptation to a changing environment. This mechanism, in comparison to others, such as cultural strength, has received less attention but it is nonetheless fundamental (exceptions are Schein, 2010; Gibbons et al, 2021c; Chatman et al, 2014).

Following Schein (2010) and Kocak and Puranam (2022) we divide culture into two dimensions: *beliefs* about the structure of the world/environment and the firm's function in it—its role, mission, means to achieve goals—and *prosocial norms*, that is, cooperative norms that prioritize help to colleagues and contributions to the whole organization (rather than focusing on one's own payoff). Using a cultural evolution model (Boyd and Richerson, 2005; Brahm and Poblete, 2021), we show that *increasing the accuracy/verisimilitude of beliefs* improves the firm's capacity to detect that a change in the environment has occurred, thereby enabling it to time innovation better and making it more productive; in turn, this provides room to expand the use of tradition and promote pro-social effort in its diffusion. This reduces the innovation versus tradition trade-off, and yields an improvement in a firm's adaptability and performance. This finding implies that, contrary to received wisdom, but consistent with some recent efforts (e.g., Gibbons et al, 2021c), cultural strength doesn't need to come at the expense of adaptation (cf., Corritore et al, 2020). Also, it connects organizational culture with the literature on cognition and organizational adaptation (Barr et al, 1992; Tripas and Gavetti, 2000; Eggers and Park, 2018) and the Carnegie School (March, 1991; Puranam, 2018; Posen et al 2018). We also find that when *prosocial norms increase its intensity*, they improve performance by enhancing the efficiency in the transmission of tradition, and this can buffer the organization from shocks. However, by promoting tradition, intense prosocial norms crowd-out innovation and thus affect the capacity of the organization to adapt; we show that if the intensity of prosocial norms is sufficiently high, this second effect dominates, and thus the relation between pro-sociality and performance is an inverted-U. We join recent efforts and calls to move beyond the view that "pro-sociality always improves performance" (Chatman et al, 2019; Hergeux et al, 2021) and study when it can have a detrimental effect on performance.

7. References

Ahuja, G., & Katila, R. 2004. Where do resources come from? The role of idiosyncratic situations. Strategic Management Journal, 25: 887-907

Akerlof, G. A., & Kranton, R. E. (2005). Identity and the Economics of Organizations. Journal of Economic perspectives, 19(1), 9-32

Argyres, N. S., & Zenger, T. R. (2012). Capabilities, transaction costs, and firm boundaries. Organization Science, 23(6), 1643-1657

Argote, L., Miron-Spektor, E. 2011 "Organizational Learning: From Experience to Knowledge." Organization Science 22 (5): 1123-1137.

Argote, L., Lee, S., & Park, J. (2021). Organizational learning processes and outcomes: Major findings and future research directions. Management science, 67(9), 5399-5429

Barnard, Ch. 1938. The functions of the executive. Harvard University Press.

Barr, P. S., Stimpert, J. L., & Huff, A. S. 1992. Cognitive change, strategic action, and organizational renewal. Strategic Management Journal, 13(S1): 15-36.

Barr, P. S. 1998. Adapting to unfamiliar environmental events: A look at the evolution of interpretation and its role in strategic change. Organization Science, 9: 644-669

Bauer, M., Blattman, C., Chytilová, J., Henrich, J., Miguel, E., & Mitts, T. (2016). Can war foster cooperation? Journal of Economic Perspectives, 30(3), 249–74.

Benner, M. J., & Tushman, M. L. (2003). Exploitation, exploration, and process management: The productivity dilemma revisited. Academy of management review, 28(2), 238-256

Besley, T. J., & Persson, T. (2022). Organizational dynamics: Culture, design, and performance. Journal of Law, Economics and Organization.

Böhm, R., and Rockenbach, B. (2013). The inter-group comparison-intra-group cooperation hypothesis: Comparisons between groups increase efficiency in public goods provision. PloS One, 8(2)e56152

Bolino, M. C., & Grant, A. M. (2016). The bright side of being prosocial at work, and the dark side, too: A review and agenda for research on other-oriented motives, behavior, and impact in organizations. Academy of Management Annals, 10(1), 599-670

Bowles, S. (2009). Did warfare among ancestral hunter-gatherers affect the evolution of human social behaviors?. Science, 324(5932), 1293-1298

Boyd, R. (2018). A different kind of animal: how culture transformed our species (Vol. 46). Princeton University Press

Boyd, R., Richerson, P.J. 1985. Culture and the Evolutionary Process. Chicago: University of Chicago Press.

Boyd, R., Richerson, P.J. 2005. The Origin and Evolution of Cultures. Oxford: Oxford University Press.

Brahm, F., Poblete, J. 2021. The evolution of productive organizations. Nature Human Behaviour

Brahm, F., & Poblete, J. (2022). Cultural Evolution Theory and Organizations. Organization Theory, 3(1), 26317877211069141

Cameron, K. S., & Quinn, R. E. (2011). Diagnosing and changing organizational culture: Based on the competing values framework. John Wiley & Sons

Cavalli-Sforza, L. L., & Feldman, M. W. (1981). Cultural transmission and evolution: A quantitative approach. Princeton University Press

Chakravarthy, B. S. 1982. Adaptation: A promising metaphor for strategic management. Academy of Management Review, 7: 35-44

Chatterji, A. K., Findley, M., Jensen, N. M., Meier, S., & Nielson, D. (2016). Field experiments in strategy research. Strategic Management Journal, 37(1), 116-132

Chatman, J. A., Caldwell, D. F., O'Reilly, C. A., & Doerr, B. (2014). Parsing organizational culture: How the norm for adaptability influences the relationship between culture consensus and financial performance in high-technology firms. Journal of Organizational Behavior, 35(6), 785-808

Chatman, J. A., & O'Reilly, C. A. (2016). Paradigm lost: Reinvigorating the study of organizational culture. Research in Organizational Behavior, 36, 199-224

Chatman, J. A., & Srivastava, S. B. (2021). A Psychologist and Sociologist Join Strategists for Breakfast: Building a Framework to Understand Culture and Its Relationship to Strategy. Strategy Science, 6(2), 119-123

Chatman, J. A., Greer, L., Sherman, E., & Doerr, B. (2019) Blurred Lines: How the Collectivism Norm Operates Through Perceived Group Diversity to Boost or Harm Group Performance in Himalayan Mountain Climbing. Organization Science 30(2):235-259

Colla, E., Ruiz-Molina, M. E., Chastenet De Gery, C., Schultz, M., Deparis, M., & Lemmet, L. (2019). Understanding franchisee performance: The role of the franchisee's autonomy, affective commitment to the network and innovativeness. International Journal of Retail & Distribution Management, 47(7), 733-751

Corritore, M., Goldberg, A., & Srivastava, S. B. (2020). Duality in diversity: How intrapersonal and interpersonal cultural heterogeneity relate to firm performance. Administrative Science Quarterly, 65(2), 359-394

Costanza, D. P., Blacksmith, N., Coats, M. R., Severt, J. B., & DeCostanza, A. H. (2016). The effect of adaptive organizational culture on long-term survival. Journal of Business and Psychology, 31(3), 361-381

Crémer, J. (1993). Corporate culture and shared knowledge. Industrial and corporate change, 2(3), 351-386

Cyert, R. M., & March, J. G. (1963). A behavioral theory of the firm (Vol. 2, No. 4, pp. 169-187)

Dada, O., Watson, A., & Kirby, D. A. (2012). Toward a model of franchisee entrepreneurship. International Small Business Journal, 30(5), 559-583

Danneels, E. (2011). Trying to become a different type of company: Dynamic capability at Smith Corona. Strategic management journal, 32(1), 1-31

Dalpiaz, E., & Di Stefano, G. (2018). A universe of stories: Mobilizing narrative practices during transformative change. Strategic Management Journal, 39(3), 664-696

Deal, T.E. & Kennedy, A.A. (1982). Corporate Cultures. The Rites and Rituals of Corporate Life. Addison-Wesley, Reading, MA

Denison, D. R., & Mishra, A. K. (1995). Toward a theory of organizational culture and effectiveness. Organization science, 6(2), 204-223

Dessein, W., & Prat, A. (2022). Organizational capital, corporate leadership, and firm dynamics. Journal of Political Economy, 130(6), 1477-1536

DiMaggio P (1994) Culture and economy. Smelser N, Swedberg R, eds. The Handbook of Economic Sociology (Princeton University Press, Princeton, NJ), 27–57

DiMaggio, P. (1997). Culture and cognition. Annual review of sociology, 23(1), 263-287

Drucker, P. F. (1994). The theory of the business. Harvard Business Review, 72(5), 95-104

Eggers, J. P., & Kaplan, S. 2009. Cognition and renewal: Comparing CEO and organizational effects on incumbent adaptation to technical change. Organization Science, 20: 461-477.

Eggers, J. P., & Kaplan, S. (2013). Cognition and capabilities: A multi-level perspective. Academy of Management Annals, 7(1), 295-340

Eggers, J. P., & Park, K. F. 2018. Incumbent adaptation to technological change: The past, present, and future of research on heterogeneous incumbent response. Academy of Management Annals, 12: 357-389

Fang, C., Lee, J., & Schilling, M. A. (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. Organization Science, 21(3), 625-642

Fehr, E. (2018). Behavioral foundations of corporate culture. University of Zurich, UBS International Center of Economics in Society, Public Paper, (7)

Felin, T., & Zenger, T. R. (2017). The theory-based view: Economic actors as theorists. Strategy Science, 2(4), 258-271

Francois, P., Fujiwara, T., & Van Ypersele, T. (2018). The origins of human prosociality: Cultural group selection in the workplace and the laboratory. Science Advances, 4(9), eaat2201

Friedkin, N. E. (1999). Choice shift and group polarization. American Sociological Review, 856-875

Gelfand, M. J., Raver, J. L., Nishii, L., Leslie, L. M., Lun, J., Lim, B. C., … & Yamaguchi, S. (2011). Differences between tight and loose cultures: A 33-nation study. science, 332(6033), 1100-1104

Gibbons, R., & Henderson, R. (2013). What do managers do?: Exploring persistent performance differences among seemingly similar enterprises. Book chapter in the Handbook of Organizational Economics, Eds. John Roberts and Robert Gibbons.

Gibbons, R., Licalzi, M., Warglien, M. 2021a. What Situation Is This? Shared Frames and Collective Performance. Strategy Science.

Gibbons, R., Siegel, J., & Weber, R. A. (2021b). Strategy meets culture (for breakfast): Understanding the relationship and highlighting its potential. Strategy Science, 6(2), 111-118

Gibbons, R. S., Grieder, M., Herz, H., & Zehnder, C. (2021c). Building an Equilibrium: Rules versus Principles in Relational Contracts, Organization Science, Articles in Advance.

Gibson, C. B., & Birkinshaw, J. (2004). The antecedents, consequences, and mediating role of organizational ambidexterity. Academy of Management Journal, 47, 209-226

Giorgi, S., Lockwood, C., & Glynn, M. A. (2015). The many faces of culture: Making sense of 30 years of research on culture in organization studies. Academy of Management Annals, 9(1), 1-54

Gorton, G. B., Grennan, J., & Zentefis, A. K. (2022). Corporate culture (No. w29322). National Bureau of Economic Research

Graham, J., Grennan, J., Campbell, H., Shivaram, R. 2022. Corporate Culture: Evidence from the Field. Journal of Financial Economics, forthcoming.

Grant, R. M. (1996). Toward a knowledge-based theory of the firm. Strategic management journal, 17(S2), 109-122

Grennan, J. (2020). Communicating culture consistently: Evidence from banks. Available at SSRN 3350645

Grennan, J. 2022. "A Corporate Culture Channel: How Increased Shareholder Governance Reduces Firm Value"

Grennan, J., & Li, K. (2022). Corporate Culture: A Review and Directions for Future Research. Available at SSRN

Groysberg, B., Lee, J., Price, J., & Cheng, J. (2018). The leader's guide to corporate culture. Harvard business review, 96(1), 44-52

Hannan, M. T., Le Mens, G., Hsu, G., Kovács, B., Negro, G., Pólos, L., ... & Sharkey, A. J. (2019). Concepts and categories: Foundations for sociological and cultural analysis. Columbia University Press

Hartnell, C. A., Ou, A. Y., & Kinicki, A. (2011). Organizational culture and organizational effectiveness: a meta-analytic investigation of the competing values framew'rk's theoretical suppositions. Journal of applied psychology, 96(4), 677

Hartnell, C. A., Ou, A. Y., Kinicki, A. J., Choi, D., & Karam, E. P. (2019). A meta-analytic test of organizational culture's association with elements of an organization's system and its relative predictive validity on organizational outcomes. Journal of Applied Psychology, 104(6), 832

Helfat, C. E., & Martin, J. A. 2015. Dynamic managerial capabilities: Review and assessment of managerial impact on strategic change. Journal of Management, 41: 1281-1312

Helfat, C. E., & Peteraf, M. A. 2015. Managerial cognitive capabilities and the microfoundations of dynamic capabilities. Strategic Management Journal, 36: 831-850

Henrich, J. 2016. The Secret of Our Success: How Culture is Driving Human Evolution, Domesticating our Species, and Making us Smarter. Princeton University Press.

Henrich, J. (2020). The weirdest people in the world: How the west became psychologically peculiar and particularly prosperous. Farrar, Straus and Giroux

Henrich, J., & Muthukrishna, M. (2021). The origins and psychology of human cooperation. Annual Review of Psychology, 72, 207-240

Hergueux, J., Henry, E., Benkler, Y., & Algan, Y. (2021). Social Exchange and the Reciprocity Roller Coaster: Evidence from the Life and Death of Virtual Teams. Organization Science (Articles in advance).

Hermalin, B. E. (2012). Leadership and corporate culture. Handbook of organizational economics, Princeton University Press 432-78

Howard-Grenville, J., Lahneman, B., & Pek, S. (2020). Organizational culture as a tool for change, Stanford Social Innovation Review, Summer.

Janis, I. L. (2008). Groupthink. IEEE Engineering Management Review, 36(1), 36

Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., & Jones, P. L. (2018). Social learning strategies: Bridge-building between fields. Trends in cognitive sciences, 22(7), 651-665

Koçak, Özgecan, Puranam, P. 2022 Decoding Culture (January 14, 2022). INSEAD Working Paper No. 2022/04/STR

Kogut, B., & Zander, U. (1996). What firms do? Coordination, identity, and learning. Organization science, 7(5), 502-518

Kreps, D. M. (1990). Corporate culture and economic theory. Perspectives on positive political economy, 90(109-110), 8

Lavie, D., Stettner, U., & Tushman, M. L. (2010). Exploration and exploitation within and across organizations. Academy of Management annals, 4(1), 109-155

Lawrence, P. and Lorsch, J 1969 Organization and Environment. Homewood, III.: Richard D. Irwin

Levinthal, D. A. (1997). Adaptation on rugged landscapes. Management science, 43(7), 934-950

Li, K., Liu, X., Mai, F., & Zhang, T. (2021). The role of corporate culture in bad times: Evidence from the COVID-19 pandemic. Journal of Financial and Quantitative Analysis, 56(7), 2545-2583

Liao, H., Su, R., Ptashnik, T., & Nielsen, J. (2022). Feeling good, doing good, and getting ahead: A meta-analytic investigation of the outcomes of prosocial motivation at work. Psychological Bulletin, 148(3-4), 158

Lix, K., Goldberg, A., Srivastava, S. B., & Valentine, M. A. (2022). Aligning differences: Discursive diversity and team performance. Management Science

March, J. G. (1991). Exploration and exploitation in organizational learning. Organization science, 2(1), 71-87

March, J. G., & Simon, H. A. (1958). Organizations. John wiley & sons

Marchetti, A. (2019). Firms of a Feather Merge Together: The Coordination Benefits of Compatible Cultures. Available at SSRN 3712838

Marchetti, A., Puranam, P. 2020 "The Cultural Attributes of 'Less Hierarchical' Forms of Organization", Working paper.

Maynard-Smith, J. (1982). Evolution and the Theory of Games. Cambridge university press

Mesoudi, A. (2011). Cultural evolution: How Darwinian theory can explain human culture and synthesize the social sciences. University of Chicago Press

Morrison, A. D., Wilhelm Jr, W. J. (2004). Partnership firms, reputation, and human capital. American Economic Review, 94(5), 1682-1692

Nowak, M. A., Sigmund, K. 2005. Evolution of indirect reciprocity. Nature, 437, 1291

O'Reilly III, C. A., & Tushman, M. L. (2013). Organizational ambidexterity: Past, present, and future. Academy of management Perspectives, 27(4), 324-338

Organ, D.W., Podsakoff, P.M. and MacKenzie, S.B., 2005.Organizational citizenship behavior:Its nature, antecedents, and consequences. Sage Publications

Posen, H. E., Keil, T., Kim, S., & Meissner, F. D. (2018). Renewing research on problemistic search—A review and research agenda. Academy of Management Annals, 12(1), 208-251

Puranam, P., Stieglitz, N., Osman, M., & Pillutla, M. M. (2015). Modelling bounded rationality in organizations: Progress and prospects. Academy of Management Annals, 9(1), 337-392

Puranam, P. (2018). The microstructure of organizations. Oxford University Press

Rand, D. G., Nowak, M. A. 2013. Human cooperation.Trends Cogn. Sci.17, 413–425

Randle, D. K., & Pisano, G. P. (2021). The evolutionary nature of breakthrough innovation: An empirical investigation of firm search strategies. Strategy Science, 6(4), 290-304

Sandvik, J. J., Saouma, R. E., Seegert, N. T., & Stanton, C. T. (2020). Workplace knowledge flows. The Quarterly Journal of Economics, 135(3), 1635-1680

Sarta, A., Durand, R., & Vergne, J. P. (2021). Organizational adaptation. Journal of management, 47(1), 43-75

Schein, E. 2010. Organizational culture and leadership. John Wiley & Sons; 4th edition
Smaldino, P. E. (2014). The cultural evolution of emergent group-level traits. Behavioral and Brain Sciences, 37(3), 243-95
Sørensen, J. B. (2002). The strength of corporate culture and the reliability of firm performance. Administrative science quarterly, 47(1), 70-91
Srivastava, S. B., & Goldberg, A. (2017). Language as a Window into Culture. California Management Review, 60(1), 56-69
Swidler, A. (1986). Culture in action: Symbols and strategies. American sociological review, 273-286
Teece, D. J. 2007. Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. Strategic Management Journal, 28: 1319-1350.
Tripsas, M., & Gavetti, G. (2000). Capabilities, cognition, and inertia: Evidence from digital imaging. Strategic management journal, 21(10-11), 1147-1161
Tripsas, M. (2009). Technology, identity, and inertia through the lens of "The Digital Photography Company". Organization science, 20(2), 441-460
Ulrich, D., Allen, J., Smallwood, N., Brockbank, W., & Younger, J. (2009). Building culture from the outside in. Strategic HR Review, 8(6), 20-27.
Uotila, J., Maula, M., Keil, T., & Zahra, S. A. (2009). Exploration, exploitation, and financial performance: Analysis of S&P 500 corporations. Strategic management journal, 30(2), 221-231
Vaara, E., Sonenshein, S., & Boje, D. (2016). Narratives as sources of stability and change in organizations: Approaches and directions for future research. Academy of Management Annals, 10(1), 495-560
Van den Steen, E. (2010a). On the origin of shared beliefs (and corporate culture). The RAND Journal of Economics, 41(4), 617-648
Van den Steen, E. (2010b). Culture clash: The costs and benefits of homogeneity. Management Science, 56(10), 1718-1738
Vicinanza, P., Goldberg, A., & Srivastava, S. (2021). Quantifying vision through language demonstrates that visionary ideas come from the periphery
Watson, A., Senyard, J., & Dada, O. L. (2020). Acts of hidden franchisee innovation and innovation adoption within franchise systems. Industrial Marketing Management, 89, 431-445
Weber RA, Camerer CF (2003) Cultural conflict and merger failure: An experimental approach. Management Science. 49(4):400–415.
Weber, K., & Dacin, M. T. (2011). The cultural construction of organizational life: Introduction to the special issue. Organization Science, 22(2), 287-298
Zenger, T. (2013). What is the theory of your firm. Harvard business review, 91(6), 72-78
Zhu, D. H. (2013). Group polarization on corporate boards: Theory and evidence on board decisions about acquisition premiums. Strategic Management Journal, 34(7), 800-822

8. Figures

Figure 1. Trade-off between innovation and help diffusing tradition. {Note: We graph the ESS equilibrium. We use c = 0.6, C = 0.8, u = 0.2, λ = 0.5, and we vary p between 0.02 and 0.20 using jumps of 0.02, to produce variance in the share of individual learners (innovation) and helping effort. Each dot of the graph is the equilibrium value of the share of individual learners and helping effort described by equations (14), (15) and (16).}

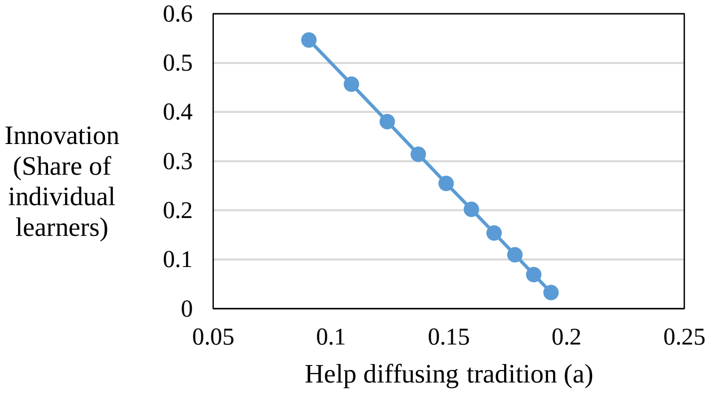

Figure 2. Impact of improving the shared beliefs. {Note: We graph the ESS equilibrium. We assume C = 0.8, c = 0.6, u = 0.2, and p = 0.1}

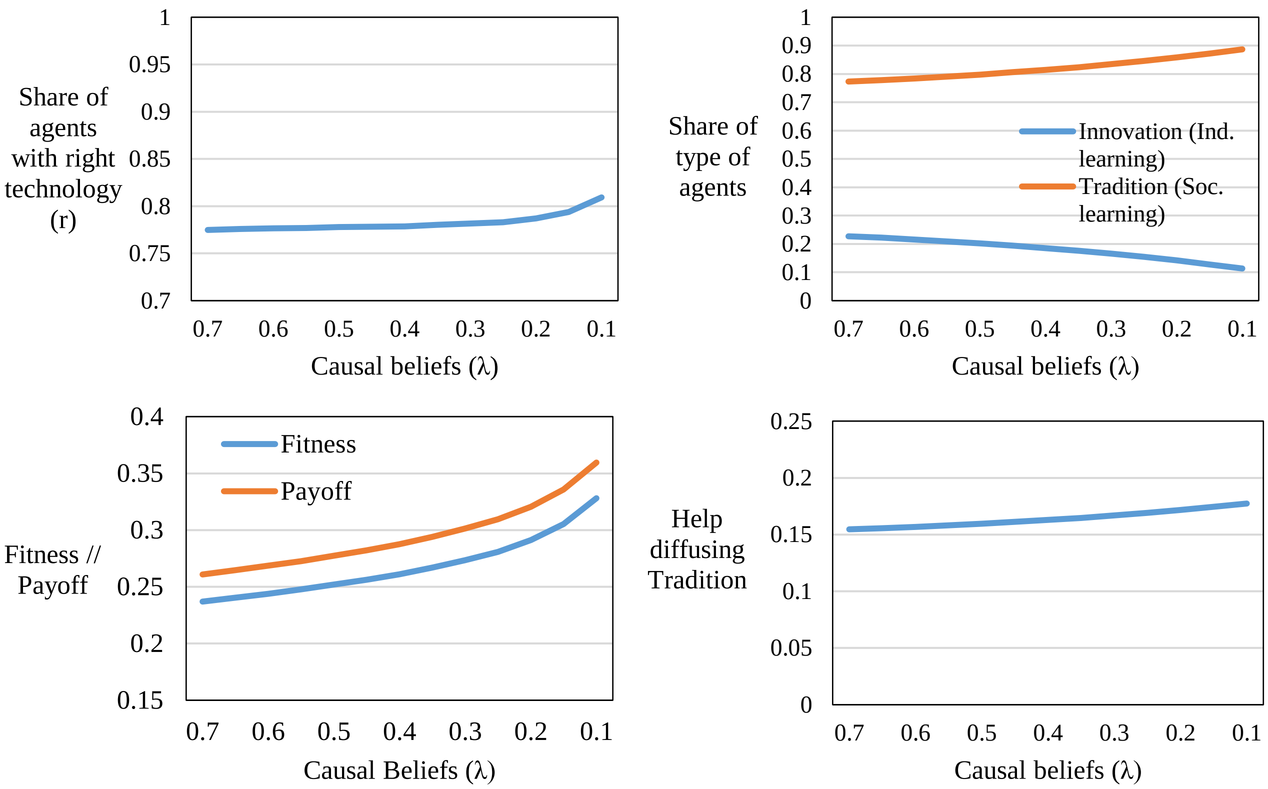

Figure 3. Innovation efficiency improves when beliefs are more accurate. {Note: We graph the ESS equilibrium. We assume $C = 0.8$, $c = 0.6$, and $p = 0.1$; in the left graph, we assume $u = 0.2$, and in the right, $\lambda = 0.5$.}

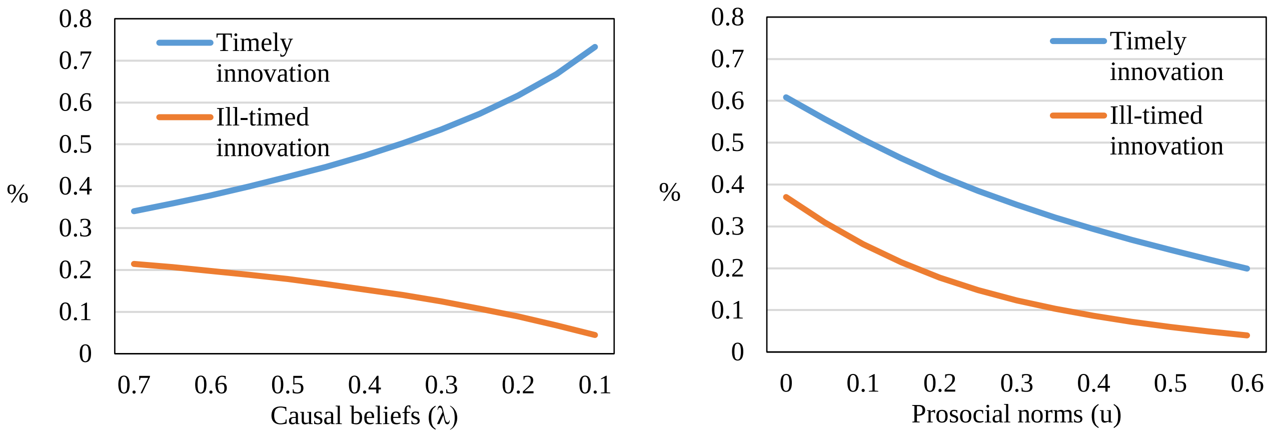

Figure 4. Improving shared beliefs reduces the trade-off between total innovation and help in diffusing tradition, and breaks it for "timely" innovation. {Note: We assume the same parameters as in the thick continuous line of Figure 2. The graph in the left uses the same information in the top right and bottom right of Figure 2. The graph of the left is equivalent to a weighted average of the graphs in the middle and right, using $p = 0.1$ and $1\text{-}p = 0.9$ as weights, respectively.}

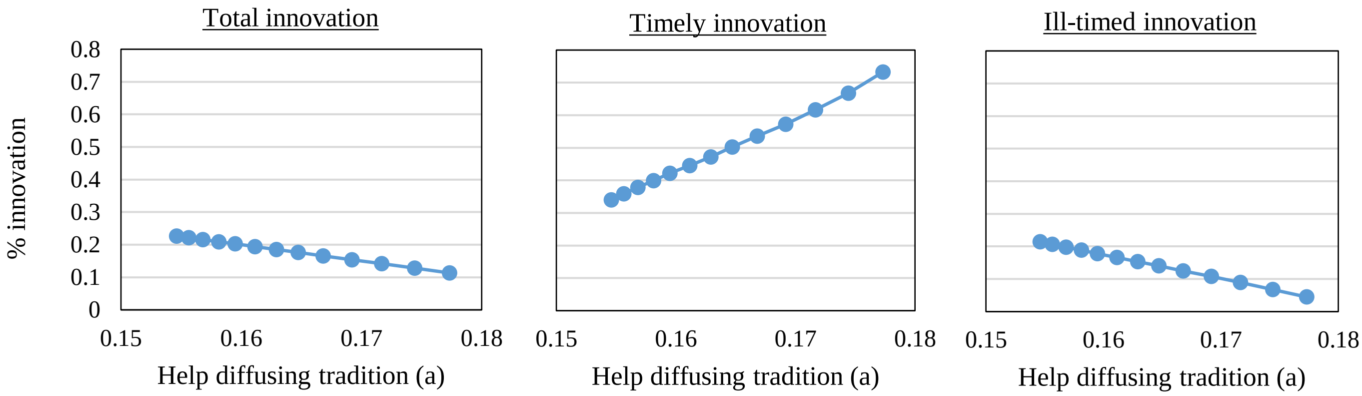

Figure 5. Dynamics of the ESS equilibrium when changing accuracy of beliefs. {Note: In all graphs, we assume $C = 0.8$, $c = 0.6$ and $u = 0.2$, and we vary the value of λ. To graph the dynamics of knowledge, we: i) use equation (6) at the equilibrium value for d, ii) seed the simulation with the equilibrium value for r in period 1, and iii) allow the world to change in period 10 (when $p = 1$) and to remain stable in the other periods (that is $p = 0$). The graph of fitness plugs the changing value of r into the fitness equation. The graphs display how the stock of adaptive knowledge and fitness behaves over time in the ESS.}

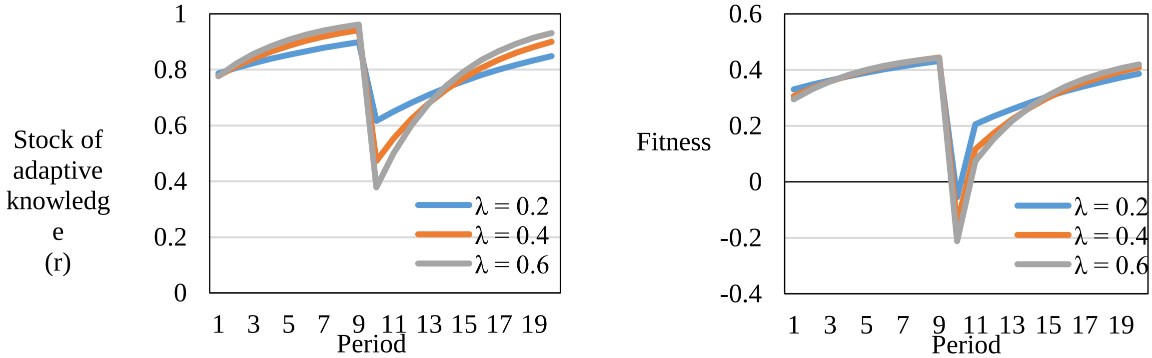

Figure 6. Impact of increasing the intensity of prosocial norms. {Note: We graph the ESS equilibrium. We assume $C = 0.8$, $c = 0.6$, and $\lambda = 0.5$. The continues thick line uses $p = 0.1$}

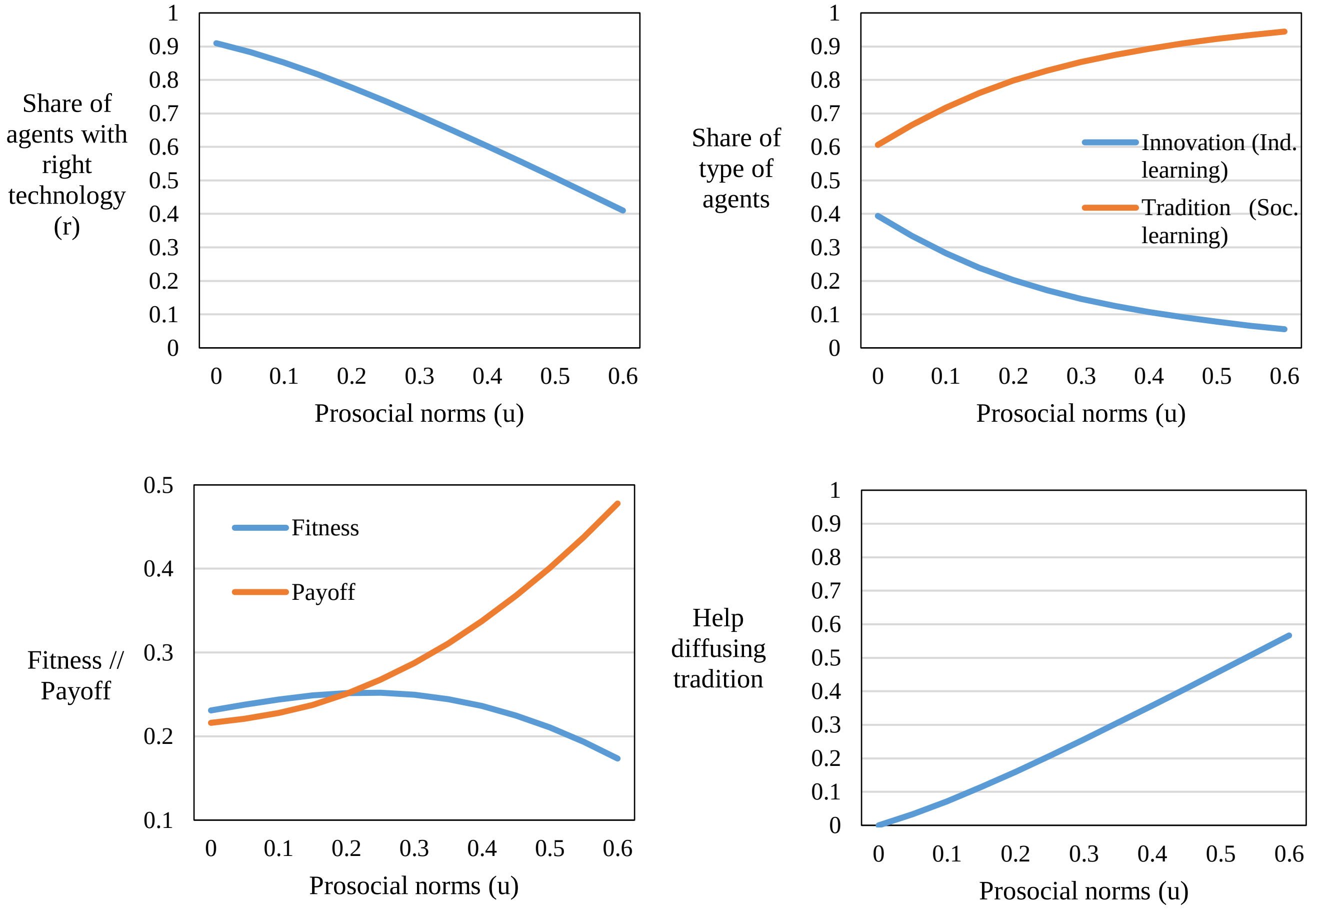

Figure 7. Dynamics of the equilibrium when changing help in diffusing tradition. {Note: In all graphs, we assume $C = 0.8$, $c = 0.6$ and $\lambda = 0.5$, and we vary the value of u.}

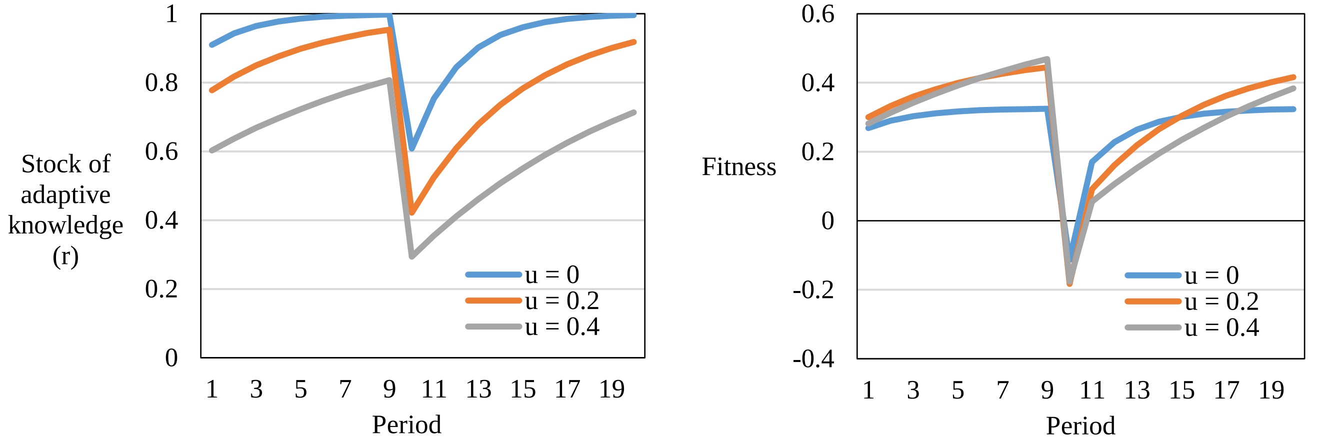

Figure 8. More intense prosocial norms enhances the trade-off between innovation and help in diffusing tradition. {Note: We assume the same parameters as in the Figure 6. The graph in the left uses the same information in the top right and bottom right of Figure 6. The graph of the left is equivalent to a weighted average of the graphs in the middle and right, using p = 0.1 and 1-p = 0.9 as weights, respectively.}

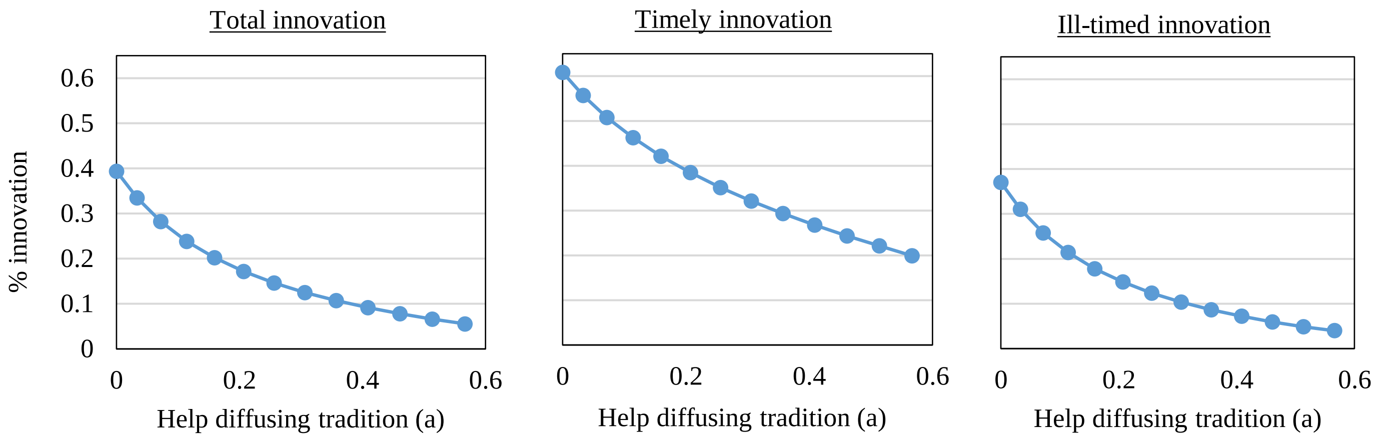

Figure 9. Complementarity between u and λ. {Note: We graph the ESS equilibrium. We assume C = 0.8, c = 0.6, and p = 0.1.}

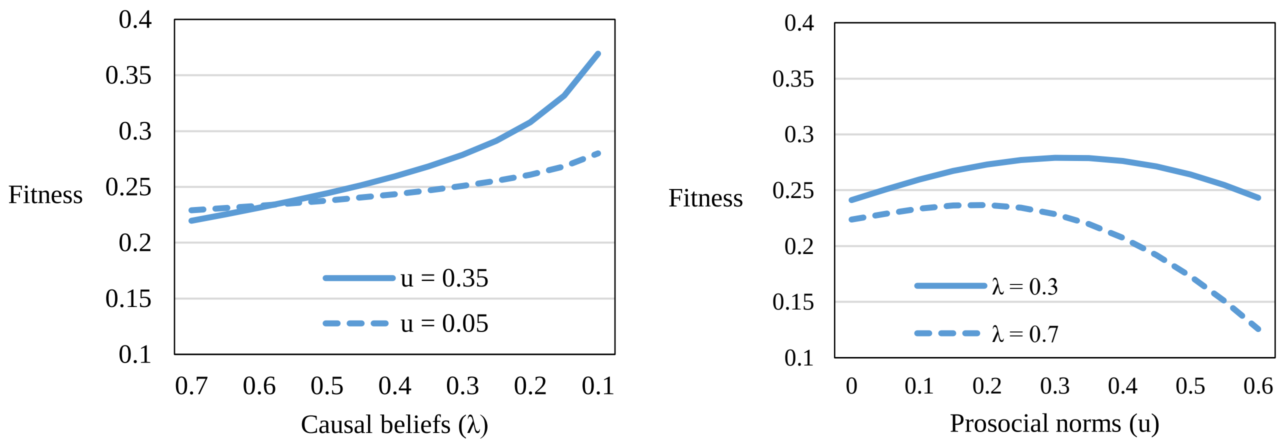

A. Online appendix

A.1. Example: Salesforce

Table A.1. How our model applies to a salesforce.

Object	Definition	Salesforce example
p	Likelihood of change in environment	Likelihood of change in customers (e.g., change in taste and/or preferences), competitors (e.g., sales strategy, complementary products/services), technology (e.g., online sales, AI), industry (e.g., entry), etc. (common to all salespersons)
r	Share of agents with tuned technology	The share of salespersons that have the right “technology/practices/skills” (e.g., sales technique, preparation before a sale, route ordering, discount management) given the state of the environment. There is a “right” technology/practices/skills because the world changes (see *p* above). This “right” technology provides higher sales level than the “wrong” technology (in our model, this distance is standardized as 1 and 0).
S	Signal obtained from the organization-environment interaction	Each period (month, quarter, semester, etc.) each salespersons receives a signal in the form of sales, customer complaints, sales returns, or non-payment (these can be absolute level or percentage change)
d	This is a threshold against which the signal is compared to decide behavior	Each salesperson has a rule of this type: “If sales drop is higher than more than 10%, then I will innovate (i.e,, try new sales techniques, reroute my day, change the way I handle discounts and promotions); if not, I will follow tradition (i.e., I will use the “technology/practices/skills” that my peers have been using in the past). The rule’s threshold varies across salespersons. Innovation is individual learning, and tradition is social learning. Salespersons learn socially either by copying directly other salespersons, usually those that are more experienced or more prestigious, but also when they learn from the organizational guidelines/handbook that have encoded accumulated experience/tradition.
λ	Beliefs about the environment and the function of the organization in it. These beliefs drive the informativeness of the signal with respect to environmental change.	How much salespersons understands: i) the environment in and of itself (e.g., customer demand, competitors, technology, industry)), ii) the function that the organizations performs for customers (or the problem it solves). These beliefs are common to all salespersons. Salesforces have different beliefs about customers: some beliefs simplistic and less accurate (e.g., “all customers are only driven by price”) which more often than not will lead to signal misinterpretation (such as not picking up that a drop in sales was caused by customers increasingly value quality); some beliefs are more accurate (e.g., “some customer prioritize price, but others prioritize quality and service”). Salesperson would also have beliefs about the function of the salesforce: some simplistic and inaccurate (e.g., “we simply take sales orders”) which will often lead to mis-interpret signals (such as sales drop driven by novel service-oriented competitors to be misinterpreted as noisy), or more sophisticated and accurate (e.g., “we provide solutions—a bundle of the right products, the right price, and a great good service—to the customer’s problem”) which will allow to pick-up a novel service-oriented competitor (a change in the environment) as a probable cause.

$\boldsymbol{\mu}$	Strength of prosocial norms in the organization	How much the organization values salespersons helping one another directly and enforces it. This translates to how much normative emphasis is placed—via informal/symbolic rewards such as reputation, awards, etc. or pro-social punishment (i.e., reprimand to selfish salespersons)—on salespersons helping their colleagues learn the sales techniques/knowledge they have (e.g., how to manage a new SKU, market information about competitors, use IT tools, a novel sales technique, etc.)
$\boldsymbol{a}$	Effort in helping others learn about tradition	Effort exerted by a salesperson so that a colleague that is imitating him/her learns better the sales techniques/knowledge he/she has (e.g., how to manage a new SKU, market information about competitors, use IT tools, etc.). In others words, effort in helping passing the tradition along to colleagues.
C, c, m(a)	Cost of learning and helping effort	Cost in terms of time and effort that the salesperson has to bear for learning and helping
$\boldsymbol{h}$	Fitness of the agent	This is the added value that the salesperson generates (i.e., price minus the "opportunity" cost C, c and m(a)). While we don't model value capture for the firm and employee (i.e., compensation plan of salesperson), we nonetheless can assume that part of this added value will be captured by the salesperson in the form of a compensation and that this share will increase with added value.
$\boldsymbol{q}$	Payoff of the agent	On top of h, salespersons care about the informal benefits their receive from helping others learn the organization's tradition. These benefits can be reputational, symbolic awards ("salesperson of the month"), or avoiding shame (e.g., the sales supervisor calling you out in public for not helping) or guilt (if the norm is internalized by the salesperson).

A.2. Threshold strategy is optimal

Remember that the expected fitness from individual learning is always $1 - C$; while from social learning is $r - c > 1 - C$ if the state does not change and 0 if it does.

Let $f(s|NC)$ and $f(s|C)$ the distributions of signal s conditional on the state not changing and the state changing respectively. Using Bayes rule we can compute the distribution of the state of the world contingent on the signal as follows:

$$P(NC|s) = \frac{f(s|NC)\,P(NC)}{f(s)}$$

$$P(C|s) = \frac{f(s|C)\,P(C)}{f(s)}$$

Conditional on the signal s, the expected payoff of individual learning is $1 - C$, which is constant in s, and the expected payoff of social learning is

$$P(NC|s)(r - c) + P(C|s)(-c)$$

For the threshold strategy to be optimal we need to show that this term is monotonically decreasing in s. Since $P(NC|s) + P(C|s) = 1$ for every s, this occurs iff $\frac{P(C|s)}{P(NC|s)}$ is increasing in s. Notice that from bayes rule $\frac{P(C|s)}{P(NC|s)} = \frac{f(s|C)}{f(s|NC)}\frac{P(C)}{P(NC)}$ which is monotonic if the two distributions $f(s|C)$ and $f(s|NC)$ satisfy the MLRP condition, which states that for any $s1 > s0$, $\frac{f(s1|C)}{f(s1|NC)} > \frac{f(s0|C)}{f(s0|NC)}$.

A.3. Proof of the existence and uniqueness of ESS

We assume that a cannot be negative and that the cost of social learning $c - a$ cannot be negative. Sufficient conditions for a to be an interior solution that satisfy these restrictions are the standard INADA conditions $\lim_{x \to c} m'(x) = \infty$ and $\lim_{x \to c} m'(0) = 0$.

Proof

We show existence and uniqueness of the ESS provided the cost of helping effort $m(a)$ is sufficiently convex. $m'(a^*) = u[(1-p)F_{\lambda=1}(d^*) + p\, F_{\lambda<1}(d^*)]$

An ESS is a triplet (a, d, r) that satisfies equations (2), (3) and (4) in the main body of the manuscript. First, notice that equation (2) can be written as

$$a = m'^{-1}(u[(1-p)F_{\lambda=1}(d) + p\, F_{\lambda<1}(d)])$$

where a is strictly increasing in d. To save notation, we define the function $a = a^*(d)$.

In an analogous way, equation (5) will be defined as $r^*(d)$. Equation $r^*(d)$ is strictly decreasing in d and $r^*(0) = 1$ and $\lim_{x \to \infty} r(x) = 0$.

Replacing both equations, the equilibrium can be characterized by the equation

$$(1-p)f_{\lambda=1}(d^*)(r^*(d) - c + a^*(d) - [1 - C]) + p\, f_{\lambda<1}(d^*)(-c + a^*(d) - [1 - C]) = 0 \qquad (6)$$

To prove existence, notice that evaluated as $d = 0$ we have $a = 0$ and $r = 1$ and, therefore, the left-hand side of equation (6) is $(1-p)(C-c) + pf_{\lambda<1}(C - c - 1)$. Because we have assumed that $C - c > p$, this term is positive.

On the other hand, as we take the limit as $d \to \infty$, we have $r = 0$ and $a \leq c$; therefore, $C - c + a - 1$ is negative and so is the left-hand side of equation (6), which is equal to $(1-p)[f_{\lambda=1}(d)](C - c + a - 1) + p[f_{\lambda<1}(d)](C - c + a - 1)$. Because function (6) is continuous, positive at d = 0 and negative as d→∞ by the intermediate value theorem, there must exist some positive value of d at which the equation becomes 0; this is an ESS.

To show uniqueness, it suffices to show that whenever equation (6) is 0, its slope is negative because in this case it crosses the axis at most at one.

Differentiating with respect to d, we obtain

$$(1-p)\left[f'_{\lambda=1}(d)\right](C-c+r^*(d)+a^*(d)-1)+p\left[f'_{\lambda<1}(d)\right](C-c+a^*(d)-1)+$$
$$(1-p)f_{\lambda=1}(d)\cdot r^{*\prime}(d)+[(1-p)f_{\lambda=1}(d)+pf_{\lambda<1}(d)]a^{*\prime}(d)$$

Considering that the function f_λ corresponds to the density of the exponential distribution, therefore $(\partial f_\lambda(d))/\partial d=-\lambda f_\lambda(d)$. Also notice that we assume that equation (6) holds, and thus the derivative can be expressed as

$$(1+\lambda)p[f_{\lambda<1}(d)](C-c+a^*(d)-1)+(1-p)f_{\lambda=1}(d)\cdot r^{*\prime}(d)+[(1-p)f_{\lambda=1}(d)+pf_{\lambda<1}(d)]a^{*\prime}(d)$$

The first term is negative because $(C-c+a^*(d)-1)$ must be negative for equation (6) to hold. The second term is also negative because $r^{*\prime}(d)$ is negative. The final term, however, is positive because $a^{*\prime}(d)$ is positive.

Finally, notice that using the implicit function theorem, we can compute

$$a^{*\prime}(d)=\frac{\mu[(1-p)f_{\lambda=1}(d)+pf_{\lambda<1}(d)]}{m''(a)}$$

This term is bounded above by $\mu/m''(a)$ and, therefore, can be made arbitrarily small if $m''(a)$ is sufficiently large. Therefore, if the function m is sufficiently convex ($m''(a)$ is sufficiently large), the ESS equilibrium is unique. ■

A.4. Impact of changing environmental uncertainty (p).

In Figure A.1 below we assume C = 0.8, c = 0.6, u = 0.2 and λ = 0.5, and then we graph the impact of modifying p between 0 and 0.2 on the model's equilibrium (p cannot be larger than 0.2, otherwise social learning never evolves given the values of C and c). Increasing p has two intuitive effects that ultimately affect fitness, displayed in the bottom-left graph of Figure A.1. First, it increases the amount of individual learning (top-right graph) which increase the stock of adaptive knowledge r (top-left graph). Second, the reduction in social learning implies that helping effort becomes less attractive, and thus it decreases (bottom-right graph). These two effects—an increase in r and a reduction in a—are slightly concave (particularly cooperation); and thus, fitness displays an inverted-U relationship with uncertainty p. Although this inverted-U implies that there is a p that maximizes fitness, the fitness reduction on each side of this maximum is approximately one-quarter of the fitness swings that are experienced when modifying other parameters (see other comparative statics below). Thus, one reading of Figure A.1 is that cultural evolution is very good at fine-tuning the parameter d and the extent of innovation vs. tradition so that fitness is not sacrificed. Finally, notice in this graph that fitness and payoff are different and may lead to tensions. Payoff is what workers care about, and in this case, they have rather less uncertainty than what the firm would prefer.

Figure A.1. Impact of changes in environmental uncertainty. {Note: We graph the ESS equilibrium. We assume C = 0.8, c = 0.6, and $\lambda = 0.5$.}

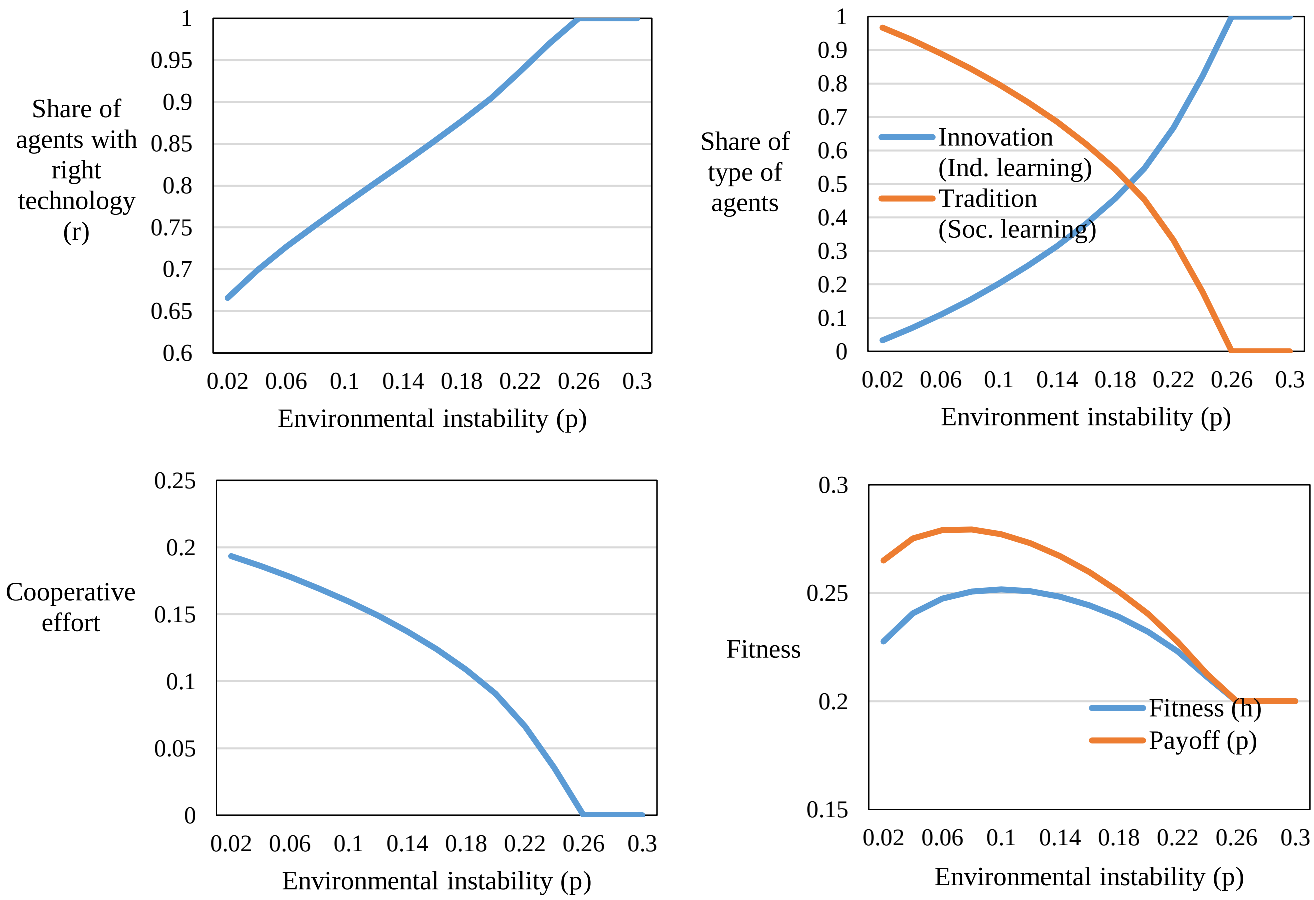

A.4.1. Condition for the evolution of social learning.

In the section 3.1 we indicated that $C - c > p$ is a necessary condition for social learning to evolve. Assuming $\lambda = 1$ when the world changed, equation (15) of the new version of the paper dictates that $f_{\lambda=1}(d^*)(-c + a^* - [1 - C]) + (1 - p)f_{\lambda=1}(d^*)\, r = 0$. Assuming that $u = 0$, and thus $a^* = 0$, this expression simplifies to $r = \frac{(1-C+c)}{(1-p)}$. Equation (16) of the new version of the paper dictates that at equilibrium $r = \frac{1-F_{\lambda=1}(d^*)}{1-(1-p)F_{\lambda=1}(d^*)}$. Therefore, $\frac{1-F_{\lambda=1}(d^*)}{1-(1-p)F_{\lambda=1}(d^*)} = \frac{(1-C+c)}{(1-p)}$. Using the exponential distribution, $\frac{(1-C+c)}{(1-p)} = \frac{\mathrm{e}^{-\mathrm{d}}}{1-(1-p)(1-\mathrm{e}^{-\mathrm{d}})}$, which after some algebra becomes $\ln\{(1-p)(1-(1-C+c)\} - \ln\{(1-C+c)\,p\} = d$. For social learning to have a positive share in equilibrium, $d > 0$, and thus $\ln\{(1-p)(1-(1-C+c)\} - \ln\{(1-C+c)\,p\} > 0$. This simplifies to $(C - c) > p$.

We just showed that the condition is necessary to have social learning in the limit case when the signal is not more informative during a period of change ($\lambda = 1$) and the prosocial norms absent ($u = 0$). Intuitively, it is not hard to grasp the necessary condition for social learning to take place will be less demanding that in this limiting case. As the signal becomes informative, agents only engage in social learning when the signal is below the threshold d, which implies that the expected fitness of social learning is larger with an informative signal. A similar effect happens as pro-social norms become stronger, which reduces the cost of social learning. Thus, the condition $(C - c) > p$ is a sufficient condition for any value $\lambda < 1$ during periods of change and $u > 0$.

A closed form analysis of the condition when $\lambda < 1$ during change and $u > 0$ is cumbersome without much added value. Instead, we revert to simulation.

The figure A.1 above shows that after a certain level of environmental uncertainty, individual learning goes to 100% and social learning goes to 0%. In this figure we use $C = 0.8$ and $c = 0.6$, and given that $\lambda < 1$ in periods of change and $u > 0$ we find that the point when social learning is selected out is close to p = 0.26, larger than the threshold 0.2 (=0.8 - 0.6), which is the value of p that grants that social learning can evolve any $\lambda < 1$ in period of change and $u > 0$.

A.5. Impact of changing learning costs C and c.

The final comparative static we perform is manipulating the costs of learning. In Figure A.2, we change the value of C from 0.7 to 0.9. The results are straightforward: if C increases, then innovation decreases and tradition increases; the stock of adaptive knowledge decreases (due to less innovation) and cooperation increases (due to more social learning). The decrease in adaptive knowledge is stronger than the increase in cooperation; thus, fitness experiences a reduction. If instead of changing C we change c, the results would be symmetric, but in opposite directions.

Figure A.2. Impact of changing the cost of individual learning. {Note: We graph the ESS equilibrium. We assume $c = 0.6$, $u = 0.2$, $p = 0.1$; and $\lambda = 0.5$}