Forthcoming American Economic Journal: Microeconomics

Learning from Others? Decision Rights, Strategic Communication, and Reputational Concerns*

Otto H. Swank and Bauke Visser

Erasmus University Rotterdam and Tinbergen Institute

September 2014

Abstract

We examine centralized versus decentralized decision-making when experience of agents is private information and communication is necessary to learn from others. An agent has reputational concerns and his market may or may not observe what the other agent chooses (global v local markets). With decentralized decision-making, agents' willingness to communicate depends heavily on what a market observes. Strikingly, less communication may improve welfare. If markets are global, centralization outperforms decentralization as it makes communication possible, and communication is informative for any finite degree of conflict among agents and with the centre.

*We thank Michael Raith and Vladimir Karamychev, and seminar audiences at Erasmus University Rotterdam, EUI, LSE, Maastricht, MEDS Kellogg School of Management, Tilburg, Bologna, Barcelona (Autònoma and Pompeu Fabra), Stockholm School of Economics, and at the MPSA conference for their comments and questions. A previous version of this paper was called Decision Making and Learning in a Globalizing World. Part of this paper was written while Visser was a Fernand Braudel fellow at the EUI.
Email: swank@ese.eur.nl and bvisser@ese.eur.nl

Keywords: centralization, decentralization, authority, learning, cheap talk, reputational concerns, policy diffusion

> *"Changing on the basis of new evidence means accepting the uncomfortable notion that we [doctors] did it wrong, or less well, before. Thus we needlessly harmed people in the past. This is painful for health professionals, (...) even if our actions were unintentional or the evidence didn't exist previously. Some find it easy to say 'Well, better stop harming now than carry on,' but denial is simpler, powerful, and comforting"*[1]

Learning from one's own experience and learning from others are two important ways in which decision-makers can improve the decisions they take over time. It may help a physician in identifying a better intervention for a patient with a given diagnosis; it may help law enforcers in fighting corporate crime more effectively; it can help organizational divisions in establishing what customer-relationship management system works best, etc. The challenge in each case is to recognize the best course of action and to ensure its diffusion.

In practice, the identification and diffusion of the best course of action raise two main problems. First, it has been established that once a decision-maker has chosen a course of action, he tends to cling to it, even if subsequently his *own experience* shows that another action would likely result in a better outcome.[2] One important reason for this conservatism has been put forward by, e.g., Kanodia, Bushman and Dickhaut (1989), and Prendergast and Stole (1996): the presence of reputational concerns. Changing course of action amounts to an admission that the previous action was inappropriate. As a result, a change affects perceptions of the ability of the decision-maker adversely. A decision-maker who wants to acquire a reputation for identifying the correct action, will be hesitant to change. The second problem is that learning from decision-makers located at other sites (hospitals, states, divisions etc.) is not automatic, but requires their willingness to share their private information. Reputational concerns may make communication strategic.

The goal of this paper is to further our understanding of learning processes by establishing how (i) the assignment of decision rights and (ii) the information on which perceptions

[1] Bewley (2008, p. 764).

[2] See, e.g., Thaler (1980).

of abilities are based jointly determine the willingness of decision-makers to share private information, the quality of the decisions taken conditional on the information transmitted, and overall welfare.

We present a simple two-period model of learning. In period one, each agent at his own site is confronted with a common problem, and receives a private signal. It may be that agents receive the same signal. The informativeness of the signal is determined by the agent's ability at identifying the better course of action. Unaware that others are struggling with the same problem, each agent optimally follows his private signal, and next privately learns the true, common value of the chosen course of action. The outcome of period one is a "historical pattern" of actions taken to address the common problem.

Next, decisions have to be made as to the action to adopt in period two. An agent may rely only on his own experience – the case studied in Prendergast and Stole. But if there is an awareness that other agents have addressed the same problem, it might be beneficial to make use of their experience. This requires communication about locally gained experience. Inspired by real world examples that we discuss in section 7.1, our analysis focuses on two dimensions that may influence the quality of learning.

(i) Second period decision rights. Do agents keep the authority to decide in period two (*decentralized decision making*), or is it in the hands of some "centre" that decides what actions are taken at the different sites (*centralized decision making*)? In the first case, communication is horizontal, among the agents. In the latter case, communication is vertical, from agents to the centre.

(ii) Information on which the perception of an agent's ability is based. As in Prendergast and Stole, we assume that (market) perceptions are based on observed actions only, not the values these actions generate, but we distinguish two cases. The perception of an agent's ability is either based on the actions taken at his site (*local markets*), or on the actions taken at all sites (*global markets*). In the latter case, comparisons across sites are possible, thanks to, e.g., increased transparency, reduced ICT costs, globalization. As highly able agents are more likely to initially take the same action than less able ones, such comparisons may affect perceptions.

We assume that the utility of an agent is increasing in the value of the action taken at his site and his end-of-period reputation, and that the centre (e.g. a health care body, the head of the police force, corporate headquarters) only cares about the value of the actions taken. We compare decentralized and centralized decision making in terms of the ex ante expected value of the actions taken in period two ('welfare'). As there may be conflicts of interest between agents, and between agents and the centre communication about the experiences gained is strategic. We focus on the case that privately gained experience constitutes unverifiable information, and that the only formal mechanism in place is the decision rights in period two. As a result, communication about the privately gained experience amounts to cheap talk.

We obtain the following main results. Under decentralized decision making, agents share information when markets are local, but do not exchange any information when markets are global. As a result, outcomes are generally better when markets are local. However, reputational concerns may distort decision making under global markets so severely that sometimes global markets perform better. Under centralized decision making, agents have incentives to misrepresent information, leading to garbled information, under both local and global markets. It follows that information exchange under centralized decision making falls between information exchange under local and global markets with decentralized decision making. Because under centralizated decision making, equilibrium decisions are not distorted, centralized decision making with global markets outperforms decentralized decision making with global markets. Finally, when reputational concerns are important, global markets lead to better outcomes than local markets under centralized decision making.

The paper is organized as follows. The next section discusses the related literature. In Section 2, we present the model. Section 3 analyses isolated agents, a benchmark situation in which agents can learn from their own past experience only. In section 4 we analyze decentralized learning, with local and global markets. In section 5 we perform the same analysis for centralized learning. Section 6 contains the comparisons. Section 7 concludes. All proofs can be found in the appendix.

1 Related Literature

Our paper is related to a number of literatures.

(1) **Information processing when information is dispersed**. Our paper is most closely related to Alonso, Dessein, and Matouschek (2008) and to Rantakari (2008). They study the desirability of a centralized or decentralized process in the context of a multidivisional firm. Each division benefits from adapting its decision to its own and privately observed market circumstances and from coordinating its decision with those of the other divisions. They show that even if coordination becomes of overriding concern to the firm, a decentralized process may still outperform a centralized process due to the difference in quality of communication.[3] As Alonso et al. and Rantakari we study the effect of the assignment of decision rights on the quality of communication and of the final decisions taken. The situation we analyse, however, is quite different. In our paper, there are no local circumstances to which a decision should ideally be adapted, nor is there a need to coordinate per se. Instead, there is room for learning from each other's past experience (to identify the better course of action), resistance to change (because of reputational concerns), and possibly the desire to convince other agents to adopt one's initial course of action (again, due to reputational concerns).

Li and Yang (2013) study cheap talk communication between two privately informed agents and a centre. Agent i is privately informed about the value of project i. Besides caring about the value of the adopted project, she is biased towards her own project. The center can implement only one of the two projects. Li and Yang (2013) study the resulting partition communication strategy and establish that simultaneous communication, sequential communication and delegation to one of the agents are outcome equivalent.[4] The key difference with our paper is that the bias in the model we study arises endogenously from reputational concerns rather than being exogenously given. Moreover, we compare centralized with decentralized decision-making.

Team theory, as developed by Marshak and Radner (1972), is one of the first formal at-

[3] Friebel and Raith (2010) study how the scope of the firm affects the quality of strategic information transmission between a division and head quarters.

[4] See Rantakari (2014) for a related model with multiplicative biases.

tempts to address the question how an organization should be structured to deal optimally with dispersed information. In this theory, the focus is on exogenously specified communication and information-processing constraints. In our paper, we focus on the effect of agents' interests on their willingness to share information. We share with the mechanism design literature a focus on the incentive problems surrounding communication. However, we do not assume that agents can commit to mechanisms. Only decision rights can be assigned. As a result, an important implication of the Revelation Principle, that a centralized process is always at least as good as a decentralized one, does not hold.[5]

There are other papers in economics and political science that explore how characteristics of decision-making processes influence the quality of cheap talk communication.[6]. The current paper differs from the existing literature in its focus on the possibilities for learning from one's own experience and from the experience of others in a context where agents have reputational concerns.

(2) **Reputational concerns**. The effect of reputational or career concerns has been studied in various environments. Holmstrom (1999) studies the incentives such concerns give to exert productive effort if there is uncertainty about an agent's ability level. If there is uncertainty about an agent's ability to 'read' or predict the state of the world one speaks of 'expert' models. Experts use the recommendations that they give, the implementation decision that they take, or the effort they exert to convince the market of their expertise.[7] Part of this expert literature looks at the effects of information disclosure ('transparency') about an expert's actions and about the outcomes of decisions.[8] The present paper is related to that literature, as the information on which an agent's reputation is based can change, either by

[5]See Poitevin (2000) and Mookherjee (2006) for excellent surveys of the assumptions underlying the Revelation Principle.

[6]In economics, see e.g. Dessein (2002, 2007), Visser and Swank (2007), Alonso, Dessein and Matouschek (2008), Rantakari (2008), and Friebel and Raith (2010). In political science, see e.g., Gilligan and Krehbiel (1987), Austen-Smith (1990), Coughlan (2000), and Austen-Smith and Feddersen (2005).

[7]Scharfstein and Stein (1990) and Ottaviani and Sorensen (2001, 2006) deal with the advice given by experts. Milbourn, Shockley and Thakor (2001) and Suurmond, Swank and Visser (2004), deal with the projects an expert implements and the effort he exerts to become informed.

[8]See Suurmond, Swank and Visser (2004) and Prat (2005) in a single-agent setting, and Levy (2007) and Swank and Visser (2009) in a committee setting.

design or by some external force, from specific to his site to involving comparisons across sites. We show that as a result of the additional information, communication is destroyed in case of a decentralized process, but improves in case of a centralized process. That is, the same form of transparency may give rise to very different effects depending on the institutions in which it is introduced.

(3) **Laboratory federalism and policy diffusion**.[9] Volden, Ting, and Carpenter (2008) study what happens if policy makers trade off policy effectiveness at solving problems and political preferences. They compare the adoption patterns of states that act independently and learn from their own past performance at addressing common problems with the patterns that arise if states learn from each other. Our focus is different from theirs as we study the quality of information exchange among decision-makers, compare centralized and decentralized decision-making, and study the effect of the informational basis of reputations.

(4) **Learning**. We already mentioned the seminal paper by Prendergast and Stole (2006) on learning from one's own observations by an agent who also cares about his perceived ability. Compared to their paper, we introduce learning from others, and hence communication, a discussion of decision rights, and different information sets on which perceptions of ability can be based. Our paper is also related to the literature on learning from others. This literature is, however, methodologically quite different from ours, as it is either *assumed* that an agent observes the value of the actions taken by others, whether the environment is strategic[10] or not[11], or that no such information is observed at all[12]. Furthermore, inertia

[9] See Oates (1999) for a survey.

[10] See the discussion of social learning in a strategic experimentation game in Bergemann and Välimäki (2006). In this literature, it is assumed that an agent perfectly observes both the technology others use and the true value they obtain. It is not clear that an agent, if he could, would not want to deviate from a strategy of truthfully revealing the value of the technology he has gained experience with. It seems that he would benefit from exaggerating the value as this would make adoption by others more likely. As a result, more (public) information would become available about this technology, and the deviator would benefit from an improved estimate of the technology's value.

[11] See Bala and Goyal (1998) for a model of learning in non-strategic networks, and Ellison and Fudenberg (1993, 1995) and Banerjee and Fudenberg (2004) for analyses of word-of-mouth communication in non-strategic environments.

[12] In the literature on informational herding, communication between decision-makers is excluded although

is an *exogenous* factor, see for example the literature on word-of-mouth communication. In our paper the quality of the information exchange and the degree of inertia are equilibrium outcomes. Were it not for the reputational concerns, the problem the agents face in our model – choosing a technology out of many – is similar to a common value bandit problem in which the bandit's arms represent the technologies of unknown, but common, value.[13] The main difference is that in a bandit problem the distribution of the value of a technology does not change with an observation of the value of *another* technology, whereas in our problem it does.

The fact that in our model the quality of information exchange and the degree of inertia are endogenous, and that a key assumption of the statistical bandit model is violated imply that a general analysis of the asymptotic behaviour of the decision-making processes described here is difficult and beyond the scope of this paper. Instead, we compare the behaviour of agents across various decision-making processes in a two-period setting.

Visser (1999) also studies the combination of learning from own past play and from the experience of others in the context of multiproduct firms that repeatedly decide what product combinations to offer. The distinction between local and global reputational markets is reminiscent of the distinction between local and global information gathering made in that paper. A firm that gathers information locally only compares its own profits with that of firms that offer overlapping product combinations, while in case of global information gathering no such limitation is assumed. Although the environment is strategic, firms cannot hide information for each other. Instead, the distinction between local and global information gathering is based on assumptions concerning information processing capacity.

Finally, our paper, by showing how reputational concerns may get in the way of learning from others and from one's own past experience, provides an explanation for why seemingly similar enterprises can show persistent performance differences. Gibbons and Henderson (2013) draw on many different sources to document this phenomenon and provide a lucid account of various hypotheses for its existence.

---

the environment in non-strategic. See e.g. Bikhchandani, Hirshleifer and Welch (1998). See Çelen, Kariv and Schotter (2010) for a first experimental analysis of social learning from actions and advice.

[13] See Bergemann and Välimäki (2006) for a concise survey of bandit problems.

(5) **Cheap talk**. In their seminal paper, Crawford and Sobel (1982) show that cheap talk between an informed sender and an uninformed receiver (decision-maker) can be informative, and that the quality of information exchange depends on the degree of alignment between the interests of both parties. In Crawford and Sobel, and in the literature on cheap talk in general, the degree of alignment is exogenously given. In our model, by contrast, it is determined in equilibrium. The reason is that senders are concerned with their reputations. These reputations are determined in equilibrium. A consequence is that, as we show below, in case of a centralized process and reputations based on comparisons across sites, cheap talk remains informative for any finite weight that agents put on their reputation.

2 The model

There are two sites (hospitals, states, etc.), $i \in \{1, 2\}$, two periods, $t = 1$ and $t = 2$, and at each site i there is an agent i. At each site and in each period, a common problem has to be addressed by using one of two technologies (policies, interventions, etc.), $X \in \{Y, Z\}$. The value of technology X is a random variable, denoted by $\widetilde{X}$, which is independent of place and time. It has a continuous and strictly increasing distribution function $F_{\widetilde{X}}(\cdot)$, and associated density function $f_{\widetilde{X}}(\cdot)$, with support $[0, 1]$. We assume that $\tilde{Y}$ and $\widetilde{Z}$ are iid, and write $F_{\widetilde{Y}} = F_{\widetilde{Z}} = F$. We use lower case letters, like x, to denote a possible value (realization) of $\tilde{X}$, such that $x \in [0, 1]$, and write $x_{i,t} = x$ to denote the realized value of technology $X_{i,t} = X$. The technology adopted at site i in period t is denoted by $X_{i,t}$.

Before $t = 1$, Nature draws y and z, and determines the ability level $\theta_i \in \{\underline{\theta}, \bar{\theta}\}$ of agent i. The ability levels and the state of the world are all statistically independent, with $\pi = \Pr(\theta_i = \bar{\theta}) \in (0, 1)$ for $i \in \{1, 2\}$.

At the beginning of period $t = 1$, both agents receive a private, unverifiable signal $s_i \in \{s^Y, s^Z\}$, $i = 1, 2$ about which technology solves the problem best. Without loss of generality, we write the analysis from the point of view of agent $i = 1$ and assume that $s_1 = s^Y$. Of course, $s_2 \in \{s^Y, s^Z\}$. The informativeness of the signal depends on the agent's ability: $\Pr(s^Y|y > z, \bar{\theta}) = 1$, $\Pr(s^Y|z > y, \bar{\theta}) = 0$, $\Pr(s^Y|y > z, \underline{\theta}) = \Pr(s^Y|z > y, \underline{\theta}) = 1/2$, for

$X \in \{Y, Z\}$.[14] That is, if i is highly able, $\theta_i = \bar{\theta}$, the signal reveals with probability one the better technology: $\Pr\left(y > z | s^Y, \bar{\theta}\right) = 1$ for $X \in \{Y, Z\}$. Hence, conditional on s^X and $\theta = \bar{\theta}$, $\widetilde{X}$ is distributed as the maximum of two iid random variables, $F_{\widetilde{X}}\left(x | s^X, \bar{\theta}\right) = F(x)^2$. On the other hand, if i is less able, $\theta_i = \underline{\theta}$, the signal is uninformative about the relative quality of the technology: $F_{\widetilde{X}}\left(x | s^X, \underline{\theta}\right) = F(x)$.[15] Note that an agent does not get a signal about his ability. Instead, π is the common prior. Still in period 1, i next decides which technology X to adopt on the basis of his signal s_i. At the end of the period he learns the value x of the chosen technology.

We distinguish three learning processes p that characterize period $t = 2$. Such a process consists of a decision-making stage, possibly preceded by a communication stage. In case there is a communication stage, agent i sends a message about the quality of the technology he adopted at site i in period $t = 1$. The receiver of this message depends on the process p. We assume that agent i, if and when he sends a message, knows the technology (*not* its value) that j has used in $t = 1$ when he sends a message. This is often the relevant case, as agents may well be aware that other technologies are used, without knowing their quality. Hence, a communication strategy $\mu_i^{\mathsf{p}}(\cdot)$ is a conditional probability distribution. Let $\mu_i^{\mathsf{p}}(m_i | s_i, x_{i,1}, X_{j,1})$ be the likelihood that i sends a cheap talk message $m_i \in M$, where $M = [0, 1]$ is a message space, in case his signal equals s_i, the observed value of $X_{i,1}$ equals $x_{i,1}$, and agent j uses technology $X_{j,1}$. Next, a decision maker determines which technology $X_{i,2}$ is adopted at site i at time $t = 2$. Who this decision maker is depends on the decision process p. Let $\mathbf{I}_i^{\mathsf{p}} \in I_i^{\mathsf{p}}$ be the information this person has at the beginning of the decision-making stage. It depends on the process p. The decision strategy d_i^{p} determines the relationship between $\mathbf{I}_i^{\mathsf{p}}$ and the technology adopted at site i at $t = 2$.

(i) In case of isolated agents (p=ia), an agent is unaware of other agents addressing the same problem, and therefore does not communicate. Hence, $I_i^{\mathsf{ia}} = \left\{s^Y, s^Z\right\} \times [0, 1]$: the information i has is his signal and the value of the technology used in $t = 1$. Agent i decides on $X_{i,2}$. Let $d_i^{\mathsf{ia}}(s_i, x_{i,1}) \in \{Y, Z\}$ denote the technology that i uses in $t = 2$ as a function of

[14] Of course, $\Pr\left(s^Z | y > z, \bar{\theta}\right) = 0$, $\Pr\left(s^Z | z > y, \bar{\theta}\right) = 1$, $\Pr\left(s^Z | y > z, \underline{\theta}\right) = \Pr\left(s^Z | z > y, \underline{\theta}\right) = 1/2$, for $X \in \{Y, Z\}$.

[15] Qualitatively, what matters for the results is that if $\theta_i = \bar{\theta}$, member i has a higher likelihood of correctly assessing the state of the economy than if $\theta_i = \underline{\theta}$.

his information.

(ii) In case of decentralized decision making (p=dl), each agent i simultaneously sends a message m_i to the other agent concerning the value of the technology he has adopted in $t = 1$. So, $I_i^{\mathsf{dl}} = \{s^Y, s^Z\} \times [0,1] \times M \times \{Y, Z\} \times M$. That is, in addition to the information in case of p=ia, and the message he sends to j, i also knows the technology $X_{j,1} \in \{Y, Z\}$ adopted at the other site, and the message $m_j \in M$ about the value of that technology. Agent i next decides on $X_{i,2}$. Let $d_i^{\mathsf{dl}}(s_i, x_{i,1}, m_i, X_{j,1}, m_j) \in \{Y, Z\}$ denote the technology that i adopts in $t = 2$ given $\mathbf{I}_i^{\mathsf{dl}}$.

(iii) In case of centralized decision making (p=cl), each agent i simultaneously sends a message m_i concerning the value of the technology he has adopted in $t = 1$ to "the centre." Hence, $I_C^{\mathsf{cl}} = \{Y, Z\}^2 \times M^2$ represents the centre's information set: information about which technology has been adopted at each site, and a message concerning the value of each technology. Next, the centre decides which technology is adopted at either site. Let $d_C^{\mathsf{cl}}(X_{1,1}, X_{2,1}, m_1, m_2) \in \{Y, Z\} \times \{Y, Z\}$ denote the function indicating for given technologies used at either site and for given messages sent by the agents the technology that is used at sites 1 and 2, respectively in $t = 2$. As no confusion can arise, we write I_C instead of I_C^{cl}, and d_C instead of d_C^{cl}.

An agent's utility depends on the value of the technology adopted at his site and on his perceived ability or reputation. This perception is based on the information set $\mathbf{\Omega}_{i,t}$. We will say that "the market" infers an agent's reputation from $\mathbf{\Omega}_{i,t}$. This market could be, e.g., the (internal) labour market or the electoral market. As in Prendergast and Stole (1996), we assume that perceptions are based on actions (technologies) chosen, not on the value generated. We distinguish two cases. Say that *markets are local* (or reputations are locally determined) if the reputation of agent i is based on the technologies used at site i only, $\mathbf{\Omega}_{i,1} = \{X_{i,1}\}$ and $\mathbf{\Omega}_{i,2} = \{X_{i,1}, X_{i,2}\}$ for $i \in \{1, 2\}$. Instead, say that *markets are global* (or reputations are globally determined) if the reputation of agent i is based on the technologies used at both sites i and j, $\mathbf{\Omega}_{i,1} = \{X_{i,1}, X_{j,1}\}$ and $\mathbf{\Omega}_{i,2} = \{X_{i,1}, X_{j,1}, X_{i,2}, X_{j,2}\}$ for $i \in \{1, 2\}$. We call $(X_{i,1}, X_{j,1}, X_{i,2}, X_{j,2})$ the *adoption vector*, indicating which technologies are adopted in $t = 1$ at sites i and j, and in $t = 2$ at sites i and j, respectively. If $X_{i,t} = X$ and the realized value is x, then the period t utility of agent i equals $x + \lambda \hat{\pi}_{i,t}(\mathbf{\Omega}_{i,t})$, where

$\hat{\pi}_{i,t}(\mathbf{\Omega}_{i,t}) = \Pr\left(\theta_i = \bar{\theta}|\mathbf{\Omega}_{i,t}\right)$ equals the belief that i is highly able conditional on $\mathbf{\Omega}_{i,t}$, and $\lambda > 0$ is the relative weight of reputational concerns. We ignore time discounting. The centre's objective is to maximize $x_{1,2} + x_{2,2}$.

It is worth emphasizing that this paper does not aim at understanding situations in which period 1 is a centrally organized experimentation stage designed to generate information on which to base period 2 decisions. Nor can decision-makers commit in $t = 1$ to ignore their private signals and decide jointly to try out all technologies possible. Instead, this paper wants to understand situations in which after some period in which experience has been gained locally, decision-makers realize that at other locations similar problems have been addressed. That is, the focus of our analysis is on period 2. Period 1 can be interpreted as history. We model history to stress that past decisions matter for current decisions, for example, through reputational concerns. Period $t = 1$ behaviour that maximizes i's utility is to follow his signal: $X_{i,1} = Y$ if and only if $s_i = s^Y$. This maximizes the expected value of the technology and minimizes the probability of changing (or having to change) technology in period 2. In section 7.2, we show that following one's signal is the unique equilibrium behaviour.

An equilibrium consists of a communication strategy $\mu_i(\cdot)$ for each agent, a belief function $f_i(\cdot|\mathbf{I})$ about the value of a technology that the decision maker did not use in $t = 1$, a decision strategy $d_i(\cdot)$ for each decision maker, and ex post reputations $\hat{\pi}_{i,t}(\cdot)$. We use the concept of Perfect Bayesian Equilibrium (from now on, equilibrium) to characterize behaviour. This requires (i) that the communication strategies are optimal for each type given decision makers' strategies and reputations; (ii) that the decision strategy is optimal given the belief functions and reputations; (iii) that beliefs and reputations are obtained using Bayes rule. We ignore babbling equilibria if an equilibrium in which information is transmitted exists.

3 Preliminary observations

3.1 Isolated agents

In this section, an agent can only learn from his own past experience. It replicates a well-known finding – reputational concerns make an agent reluctant to correct past decisions. The purpose is to provide a benchmark against which the welfare gains stemming from the possibility to learn from the experience of others can be measured.

Once agent 1 has followed his signal s^Y in period 1 and observed value y, he has to decide whether to continue with his technology. Note that having received s^Y and next observing y allows an agent to update the expected value of the other technology,

$$E\left[\widetilde{Z}|s^Y,y\right]=\Pr\left(\bar{\theta}|s^Y,y\right)E\left[\widetilde{Z}|s^Y,y,\bar{\theta}\right]+\Pr\left(\underline{\theta}|s^Y,y\right)E\left[\widetilde{Z}\right], \tag{1}$$

where we have used that $E\left[\widetilde{Z}|s^Y,y,\underline{\theta}\right]=E\left[\widetilde{Z}\right]$. Two effects of y can be distinguished. First, the larger is y, the more likely it is that agent 1 is highly able and correctly identified the more valuable technology. This is the $\Pr\left(\bar{\theta}|s^Y,y\right)$ term. Second, conditional on the agent being highly able, a higher value of y increases the expected value of $\widetilde{Z}$. This is the $E\left[\widetilde{Z}|s^Y,y,\bar{\theta}\right]$ term. Of course, $E\left[\widetilde{Z}|s^Y,y,\bar{\theta}\right]\leq E\left[\widetilde{Z}\right]$ for all y. The following lemma summarizes some characteristics of $E\left[\widetilde{Z}|s^Y,y\right]$.

Lemma 1 *The expected value of $\widetilde{Z}$ given $s_i=s^Y$ and y satisfies: (a) $E\left[\widetilde{Z}|s^Y,0\right]=E\left[\widetilde{Z}|s^Y,1\right]=E\left[\widetilde{Z}\right]$, and $E\left[\widetilde{Z}|s^Y,y\right]<E\left[\widetilde{Z}\right]$ for $y\in(0,1)$; (b) $E\left[\widetilde{Z}|s^Y,y\right]$ is decreasing in y for $y<E\left[\widetilde{Z}|s^Y,y\right]$, increasing for $y>E\left[\widetilde{Z}|s^Y,y\right]$, and $y=E\left[\widetilde{Z}|s^Y,y\right]$ has a unique solution.*

This lemma is illustrated in Figure 1, panel a. The dashed (horizontal) line represents the unconditional expectation $E\left[\widetilde{Z}\right]$, and the drawn line the conditional expectation $E\left[\widetilde{Z}|s^Y,y\right]$.

Ignore reputational concerns for the moment. Given $\mathbf{I}_1^{\mathsf{ia}}=\left\{s^Y,y\right\}$, the decision strategy that maximizes the expected value of the technology adopted at site 1 in the second period, the first-best strategy, is to stick to the existing technology if and only if $y\geq E\left[\widetilde{Z}|s^Y,y\right]$.

Figure 1: Isolated Agents. Panel (a) depicts the first-best threshold value and the equilibrium threshold value $\bar{y}_{\mathsf{ia}}$ for $\lambda < \bar{\lambda}_{\mathsf{ia}}$; panel (b) reports the equilibrium values for $f_X = 1$ and $\pi = 1/2$. Thus, $\bar{\lambda}_{\mathsf{ia}} = 1$. Note that $\Delta\hat{\pi}^*$ is the equilibrium reputational gap.

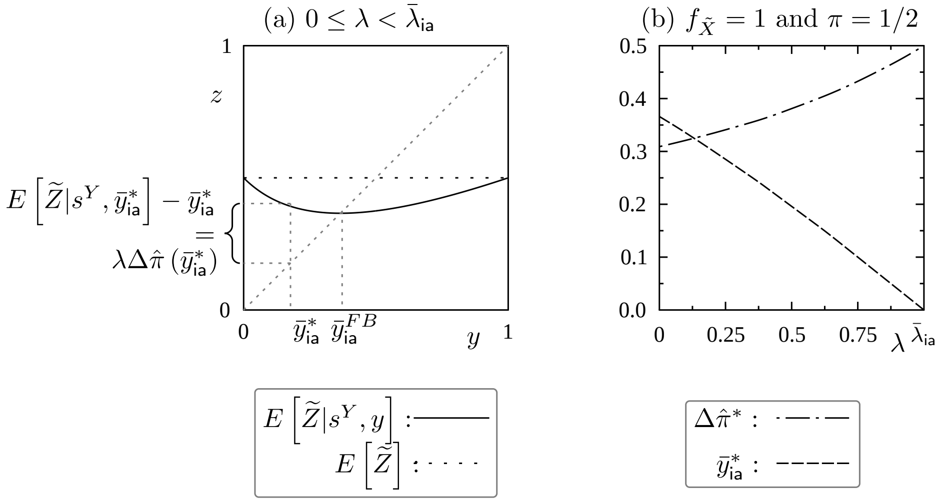

It follows from lemma 1, part (b), and it is clear from Figure 1, panel (a), that the first-best decision strategy is a *single-threshold strategy*,

$$d_1^{\mathsf{ia}}\left(\mathbf{I}^{\mathsf{ia}}; \bar{t}\right) = \begin{cases} Y & \text{if } y \geq \bar{y} \\ Z & \text{otherwise,} \end{cases}$$

with $\bar{y} = \bar{y}^{FB}_{\mathsf{ia}}$ and where $\bar{y}^{FB}_{\mathsf{ia}}$ solves $\bar{y}^{FB}_{\mathsf{ia}} = E\left[\widetilde{Z}|s^Y, \bar{y}^{FB}_{\mathsf{ia}}\right]$.

Besides being interested in picking the most valuable technology, an agent is also interested in his reputation. Consider a single-threshold decision strategy and any threshold value $\bar{t} \in (0,1)$. In case of isolated agents, markets only have local knowledge. Let $\hat{\pi}\left(Y, X_{1,2}; \bar{y}\right)$ denote the reputation, obtained using Bayes' rule, for $X_{1,2} \in \{Y, Z\}$, and if the agent uses the threshold $\bar{y}$. Then,[16]

$$\hat{\pi}_1\left(Y, Y; \bar{y}\right) = \frac{1 + F\left(\bar{y}\right)}{1 + F\left(\bar{y}\right)\pi}\pi > \pi > \hat{\pi}_1\left(Y, Z; \bar{y}\right) = \frac{F\left(\bar{y}\right)}{F\left(\bar{y}\right)\pi + (1-\pi)}\pi. \tag{2}$$

Irrespective of $\bar{y}$, continuation commands a higher reputation than switching to the other technology. Continuation suggests having observed a sufficiently high value of y. A highly able agent is more likely to have implemented a technology that generates a high value than

[16]Deriviations can be found in the proof of Proposition 1 in the Appendix.

a less able agent. Hence, as an agent cares about his reputation, he wants to deviate from the first-best decision rule by lowering the hurdle that his initial technology should pass for its continuation. The agent wants to give up technological adequacy for reputational benefits. We will call the difference $\hat{\pi}_1(Y,Y;\bar{y}) - \hat{\pi}_1(Y,Z;\bar{y})$ the *reputational gap*. It is the source of the distortion. Proposition 1 describes equilibrium behaviour of an isolated agent.

Proposition 1 *In case of* isolated agents, *and for* $\lambda < \bar{\lambda}_{\mathsf{ia}} = E\left[\widetilde{Z}\right]/\pi$, *there exists an equilibrium in which the decision strategy is a single-threshold strategy with threshold value* $\bar{y}^{\mathsf{ia}}$ *that satisfies*

$$\lambda\left[\hat{\pi}_1\left(Y,Y;\bar{y}^{\mathsf{ia}}\right) - \hat{\pi}_1\left(Y,Z;\bar{y}^{\mathsf{ia}}\right)\right] = E\left[\widetilde{Z}|s^Y,\bar{y}^{\mathsf{ia}}\right] - \bar{y}^{\mathsf{ia}}, \tag{3}$$

with $\bar{y}^{\mathsf{ia}} \in \left(0, \bar{y}_{\mathsf{ia}}^{FB}\right)$. $\bar{y}^{\mathsf{ia}}$ *is a decreasing function of* λ.[17] *For* $\lambda \geq \bar{\lambda}_{\mathsf{ia}}$, $\bar{y}^{\mathsf{ia}} = 0$, *i.e., agent* 1 *always continues his initial technology, so* $\hat{\pi}_1(Y,Y;0) = \pi$. *A plausible out-of-equilibrium is* $\hat{\pi}_1(Y,Z;0) = 0$.

Eq (3) is illustrated in Figure 1, panel (b). At the threshold value $\bar{y}^{\mathsf{ia}}$ the agent is indifferent between sticking to Y and switching to Z. It follows from (2) that if the hurdle for continuation is lowered, passing the hurdle becomes a less convincing signal of ability. At the same time, not passing a lower hurdle becomes a stronger signal of incompetence. Figure 1, panel (c) illustrates the proposition for a uniform distribution and $\pi = 1/2$. It shows the equilibrium values of $\bar{y}^{\mathsf{ia}}$ and $\hat{\pi}_1\left(Y,Y;\bar{y}^{\mathsf{ia}}\right) - \hat{\pi}_1\left(Y,Z;\bar{y}^{\mathsf{ia}}\right)$, denoted by $\Delta\hat{\pi}^*$ to save space.

3.2 First-best behavior never part of equilibrium

In the rest of the paper, agents potentially can learn from the experience of other agents. From a project value perspective, it would be best if each agent truthfully reveals the value of the technology he adopted in period 1 such that period-two decision makers are well-informed. Say that 1 *truthfully reveals* his private information if, for all $y \in [0,1]$, and all $X_{2,1} \in \{Y,Z\}$, $\Pr\left(m_1|s^Y,y,X_{2,1}\right) = 1$ if $m_1 = y$ and $\Pr\left(m_1|s^Y,y,X_{2,1}\right) = 0$ otherwise.

[17] We cannot exclude the possibility of multiple equilibria in general. In case of multiple equilibria, we show that the highest and the lowest equilibrium values of $\bar{y}^{\mathsf{ia}}$ are decreasing functions of λ. We have established numerically that in case of the uniform distribution, the equilibrium is unique.

Next, in period two decision makers should adopt the technology with the highest expected value conditional on the information obtained. We call such behaviour first-best behaviour. Before turning to a detailed analysis of decentralized and centralized decision making in case of local or global markets, we argue here, and formally prove in Appendix A, that first-best behaviour is not equilibrium behaviour in any of these four cases. The main reason is as in the case of isolated agents. If agents were to use first-best behaviour, continuation of a technology would command a higher reputation than changing technology. Essentially, in case of decentralized decision-making, and assuming the other agent sticks to first-best behaviour, an agent is willing to hold on to his technology for values that are just smaller than the value reported by the other agent to enjoy the full reputational benefits at the cost of only a small loss in project value in period two. If instead the center decides about technology adoption in period two, and assuming again that the other agent tells the truth, then an agent can slightly exaggerate the information he sends to the center in a bid to enjoy the full reputational benefits of having his technology continued at the cost of only a small loss in project value.

4 Decentralized decision making

In this section, the right to decide about the technology to be adopted in period two remains with the agents. Before an agent communicates with the other agent, he learns whether the other agent used the same or the other technology in $t = 1$. With global reputations, the market of an agent also observes what technology the other agent used. Thus, there are two reputational gaps, one determining the size of the distortion for the situation in which both agents used the same technology in $t = 1$, and one for the situation where they used different technologies. If instead reputations are locally determined, by definition the market of an agent is not able to see what technology the other agent used. As a result, there is only one reputational gap, and the distortion is equally large for these two situations. In the next subsections, we show how these reputations are formed and how sensitive they are to switching technology. In 4.1 we study equilibrium behaviour with local markets, and in 4.2 we turn to global markets. In 6.1, we compare the performance of decentralized learning under both types of reputation formation.

In case of decentralized decision-making, the decision strategy is a *double-threshold strategy* $d_1^{\mathsf{dl}}\left(\mathbf{I}_1^{\mathsf{dl}};\bar{y}_{\mathsf{S}},\bar{y}_{\mathsf{D}}\left(E\left[\widetilde{Z}|\mathbf{I}_1^{\mathsf{dl}}\right]\right)\right)$ with thresholds $\left(\bar{y}_{\mathsf{S}},\bar{y}_{\mathsf{D}}\left(E\left[\widetilde{Z}|\mathbf{I}_1^{\mathsf{dl}}\right]\right)\right)\geqq 0$,

$$d_1^{\mathsf{dl}}\left(\mathbf{I}_1^{\mathsf{dl}};\bar{t}_{\mathsf{S}},\bar{t}_{\mathsf{D}}\right)=\begin{cases} Y & \text{if } X_{2,1}=Y \text{ and } y\geq\bar{y}_{\mathsf{S}} \\ Y & \text{if } X_{2,1}=Z \text{ and } y\geq\bar{y}_{\mathsf{D}}\left(E\left[\widetilde{Z}|\mathbf{I}_1^{\mathsf{dl}}\right]\right) \\ Z & \text{otherwise.}\end{cases} \tag{4}$$

That is, agent 1 continues with his original technology Y (i) if both agents used the same technology and its value exceeds $\bar{y}_{\mathsf{S}}$; or (ii) if the agents used different technologies, and the current technology is better than $\bar{y}_{\mathsf{D}}\left(E\left[\widetilde{Z}|\mathbf{I}_1^{\mathsf{dl}}\right]\right)$. In the latter case, the threshold value can depend on the information that 1 obtains from the other agent.

4.1 Local markets

With local markets, agent 1's reputation is independent of what agent 2 decides. As a result, his communication strategy is payoff irrelevant, and any such strategy can be sustained in equilibrium. Therefore, if they can coordinate their communication strategies, they would coordinate on the welfare-maximizing strategy. Here we proceed by assuming that coordination is not possible. Then, if we assume even an infinitesimal cost of lying (see e.g. Kartik 2009), the agent will choose to report truthfully. Absent any motive to influence the other agent, the quality of information exchange is high. Once communication has taken place, each agent independently decides whether to continue with his original technology or to switch to the other technology using the decision strategy in (4), with $E\left[\widetilde{Z}|\mathbf{I}_1^{\mathsf{dl}}\right]=z$. First-best threshold values would equal $(\bar{y}_{\mathsf{S}},\bar{y}_{\mathsf{D}}(z))=\left(\bar{y}_{\mathsf{dl}}^{FB},z\right)$, where $\bar{y}_{\mathsf{S}}^{FB}$ satisfies $\bar{y}_{\mathsf{S}}^{FB}=E\left[\widetilde{Z}|s^Y,s^Y,\bar{y}_{\mathsf{S}}^{FB}\right]$.[18] Because of reputational concerns in equilibrium $\bar{y}_{\mathsf{S}}<\bar{y}_{\mathsf{dl}}^{FB}$ and $\bar{y}_{\mathsf{D}}(z)<z$ for all z. As the difference between first-best threshold value and the equilibrium threshold values equals λ times the reputational gap, we can rewrite $\bar{y}_{\mathsf{D}}(z)=z-\bar{t}_{\mathsf{D}}$. Thus, if both agents adopted Y in $t=1$, then agent 1 sticks to this technology if and only if $y+\lambda\hat{\pi}_1(Y,Y;\bar{y}_{\mathsf{S}},\bar{t}_{\mathsf{D}})\geq E\left[\widetilde{Z}|s^Y,s^Y,y\right]+\lambda\hat{\pi}_1(Y,Z;\bar{y}_{\mathsf{S}},\bar{t}_{\mathsf{D}})$. Sim-

[18] Of course, the fact that both experts used the same technology in the first period bodes well for the superiority of this technology: $\bar{y}_{\mathsf{S}}^{FB}<\bar{y}_{\mathsf{ia}}^{FB}$.

ilarly, in case agents adopted different technologies, agent 1 wants to continue with Y iff $y+\lambda\hat{\pi}_1(Y,Y;\bar{y}_{\mathsf{S}},\bar{t}_{\mathsf{D}}) \geq z+\lambda\hat{\pi}_1(Y,Z;\bar{y}_{\mathsf{S}},\bar{t}_{\mathsf{D}})$. Proposition 2 describes equilibrium behaviour. Note that lo stands for local markets.

Proposition 2 *Define* $\underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}} = E\left[\widetilde{Z}\right]/\hat{\pi}_1\left(Y,Y;0,E\left[\widetilde{Z}\right]\right)$ *and* $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}} = 1/\pi$. *In case of decentralized decision making and local markets, in equilibrium*
(i) truthful revelation is the communication strategy;
(ii) the belief functions are $\Pr(x_{2,1}|m_2) = 1$ *for* $x_{2,1} = m_2$ *and* $\Pr(x_{2,1}|m_2) = 0$ *for* $x_{2,1} \neq m_2$*;*
(iii) the decision strategy is a double-threshold strategy. For $\lambda < \underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$*, threshold values* $\left(\bar{y}_{\mathsf{S}}^{\mathsf{lo}},\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right)$ *satisfy*

$$\lambda\left[\hat{\pi}_1\left(Y,Y;\bar{y}_{\mathsf{S}}^{\mathsf{lo}},\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right)-\hat{\pi}_1\left(Y,Z;\bar{y}_{\mathsf{S}}^{\mathsf{lo}},\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right)\right] = E\left[\widetilde{Z}|s^Y,s^Y,\bar{y}_{\mathsf{S}}^{\mathsf{lo}}\right]-\bar{y}_{\mathsf{S}}^{\mathsf{lo}} \tag{5}$$

$$\lambda\left[\hat{\pi}_1\left(Y,Y;\bar{y}_{\mathsf{S}}^{\mathsf{lo}},\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right)-\hat{\pi}_1\left(Y,Z;\bar{y}_{\mathsf{S}}^{\mathsf{lo}},\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right)\right] = \bar{t}_{\mathsf{D}}^{\mathsf{lo}}, \tag{6}$$

with $\bar{y}_{\mathsf{S}}^{\mathsf{lo}} \in \left(0,\bar{y}_{\mathsf{S}}^{FB}\right)$ *and* $\bar{t}_{\mathsf{D}}^{\mathsf{lo}} \in (0,1)$. *For* $\lambda \in [\underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}},\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}})$*, threshold values are* $\left(0,\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right)$ *and* $\bar{t}_{\mathsf{D}}^{\mathsf{lo}}$ *solves* $\lambda\hat{\pi}_1\left(Y,Y;0,\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right) = \bar{t}_{\mathsf{D}}^{\mathsf{lo}}$. *Finally, for* $\lambda \geq \bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$*, threshold values equal* $(0,1)$.

Figure 2, panels (a) and (b) show the structure of the equilibrium. For $\lambda < \underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$, see panel (a) and Eqs (5) and (6), the size of the distortions, $E\left[\widetilde{Z}|s^Y,s^Y,\bar{y}_{\mathsf{S}}^{\mathsf{lo}}\right]-\bar{y}_{\mathsf{S}}^{\mathsf{lo}}$ and $\bar{t}_{\mathsf{D}}^{\mathsf{lo}}$, and the value of the reputational gap, $\lambda\left[\hat{\pi}_1\left(Y,Y;\bar{y}_{\mathsf{S}}^{\mathsf{lo}},\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right)-\hat{\pi}_1\left(Y,Z;\bar{y}_{\mathsf{S}}^{\mathsf{lo}},\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right)\right]$, are the same. The loss in technological value due to the distortion should in either case be compensated by the same boost in reputation.

For $\lambda \in [\underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}},\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}})$, illustrated in panel (b), if 1 learns that 2 used the *same* technology, he continues his initial technology irrespective of its value y, $\bar{y}_{\mathsf{S}}^{\mathsf{lo}} = 0$, whereas if 1 learns that 2 used a *different* technology, 1 may still change technology. For $\lambda \geq \bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$, 1 sticks to his initial technology Y, irrespective of its value y, and regardless of what 2 reports, $\left(\bar{y}_{\mathsf{S}}^{\mathsf{lo}},\bar{t}_{\mathsf{D}}^{\mathsf{lo}}\right) = (0,1)$. Then $\hat{\pi}_1(Y,Y;0,1) = \pi$ as continuation of Y does not reveal any information on ability, while $\hat{\pi}_1(Y,Z;0,1) = 0$ is a plausible out-of-equilibrium belief. Hence, $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}} = 1/\pi$. Panel (c) illustrates the reputational gap and the threshold values for the uniform distribution and $\pi = 1/2$.

As discussed above, with locally determined reputations an agent's communication strategy is payoff irrelevant. So far in this section we have focused on truth telling. The

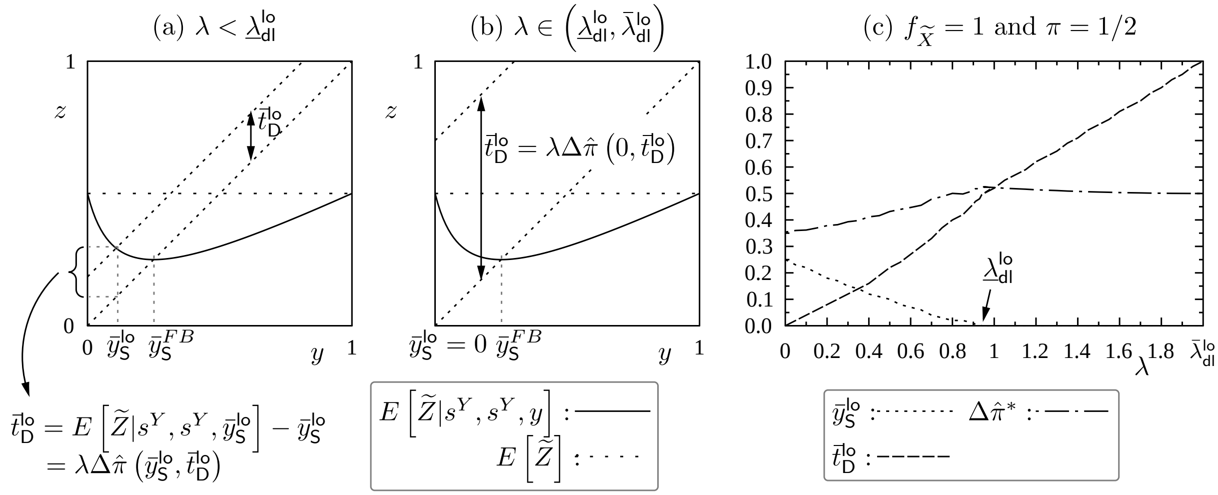

Figure 2: Decentralized decision making and local markets. Panels (a) and (b) depict the structure of equilibrium. Panel (c) reports equilibrium threshold values and the reputational gap for the uniform distribution and $\pi = 1/2$. Hence, $\underline{\lambda}^{\mathsf{lo}}_{\mathsf{dl}} < 1$ and $\bar{\lambda}^{\mathsf{lo}}_{\mathsf{dl}} = 2$.

essence of what happens in case agents do not truthfully share their information can best be seen by assuming that agents do not share any information at all. Let us first examine the implication of this alternative assumption for agents' technology decisions in case they used different technologies in period 1. Then, agent 1 continues with Y in period 2 if $y + \lambda\hat{\pi}_1\left(Y, Y; \bar{y}_{\mathsf{S}}, \bar{t}_{\mathsf{D}}\right) \geq E\left[\tilde{Z}|s^Y, s^Z, y\right] + \lambda\hat{\pi}_1\left(Y, Z; \bar{y}_{\mathsf{S}}, \bar{t}_{\mathsf{D}}\right)$. If for $y = 0$, agent 1 continues with Y, $\hat{\pi}_1\left(Y, Z; 0, \bar{t}_{\mathsf{D}}\right) = 0$ is a plausible out-of equilibrium belief. Moreover, $E\left[\tilde{Z}|s^Y, s^Z, 0\right] = E\left[\tilde{Z}\right]$. The highest value of λ for which technology switching occurs equals $\bar{\lambda}^{\mathsf{lo}}_{\mathsf{dl}} = E\left[\tilde{Z}\right]/\pi < 1/\pi$. No information sharing therefore narrows the range of λ for which an agent may switch technology in case they used different technologies in period 1. The reason is that when information is shared, agent 1 can learn that $z = 1$. Such a high value tremendously discourages agent 1 to continue with a bad technology Y.

Now suppose that both agents used technology Y in period 1. Then, not sharing information in case they used different technologies in period 1 affects agents' behaviour through the posteriors $\hat{\pi}_i\left(Y, Y; \bar{y}^{\mathsf{lo}}_{\mathsf{S}}, \bar{t}^{\mathsf{lo}}_{\mathsf{D}}\right)$ and $\hat{\pi}_i\left(Y, Z; \bar{y}^{\mathsf{lo}}_{\mathsf{S}}, \bar{t}^{\mathsf{lo}}_{\mathsf{D}}\right)$. When information is shared at $\lambda = \underline{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}$, switching technology would destroy an agent's reputation as it implies $z > y$: $\hat{\pi}_i\left(Y, Z; \bar{y}^{\mathsf{lo}}_{\mathsf{S}}, \bar{t}^{\mathsf{lo}}_{\mathsf{D}}\right) = 0$. As a result the reputational gap is wide. When information is not shared and $\lambda = \underline{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}$, switching does not necessarily imply that $z > y$, hence $\hat{\pi}_i\left(Y, Z; \bar{y}^{\mathsf{lo}}_{\mathsf{S}}, \bar{t}^{\mathsf{lo}}_{\mathsf{D}}\right) >$

0. This narrows the reputational gap, meaning that reputational concerns provide weaker incentives to distort the technology decision. The implication is that not sharing information reduces distortions in technology decisions for moderate values of λ in case agents used the same technology. For this reason, we cannot exclude that for moderate values of λ not sharing information increases welfare. For large values of λ, however, sharing information increases welfare.

4.2 Global markets

The key difference between local and global markets is that in the latter case, agent 1's ex post reputation is strengthened by a switch by agent 2 to 1's initial technology, $\hat{\pi}_1(Y, Z, Y, Z) < \hat{\pi}_1(Y, Z, Y, Y)$ and $\hat{\pi}_1(Y, Z, Z, Z) < \hat{\pi}_1(Y, Z, Z, Y)$. As an agent wants to convince the other agent to switch technology, all meaningful communication between the two agents is destroyed. The unique equilibrium communication strategy in case $X_{2,1} = Z$ is a pooling strategy.[19] If instead agents initially adopted the *same* technology, Y, it is easy to see that truthful revelation is an equilibrium strategy. Communication is also irrelevant.[20] Once communication has taken place, each agent decides whether to continue with his original technology or to switch to the other technology using the decision strategy in (4), with $E\left[\widetilde{Z}|\mathbf{I}_1^{\mathsf{dl}}\right] = E\left[\widetilde{Z}\right]$. Thus, $\bar{y}_{\mathsf{D}}\left(E\left[\widetilde{Z}|\mathbf{I}_1^{\mathsf{dl}}\right]\right)$ reduces to $\bar{y}_{\mathsf{D}}$. Of course, conditional on the information exchanged, the values of $\bar{y}_{\mathsf{S}}$ and $\bar{y}_{\mathsf{D}}$ that would maximize the expected technological value in $t = 2$ are $\bar{y}_{\mathsf{S}} = \bar{y}_{\mathsf{S}}^{FB}$, and $\bar{y}_{\mathsf{D}} = E\left[\widetilde{Z}\right]$. The next proposition describes equilibrium behaviour. Note that gl stands for globally determined reputations. [21]

Proposition 3 *Define* $\underline{\lambda}_{\mathsf{dl}}^{\mathsf{gl}} = E\left[\widetilde{Z}\right]\frac{1+\pi^2}{\pi(1+\pi)}$ *and* $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{gl}} = E\left[\widetilde{Z}\right]\frac{1+\pi}{\pi}$*. In case of decentralized decision making with global markets, there exists a unique equilibrium in which*

(i) the communication strategy is (a) a pooling strategy if initial technologies differ, and (b)

[19] To avoid a discussion of out-of-equilibrium beliefs, we assume that each agent uses a probability distribution over the full support $[0, 1]$ that is independent of the value x he observed. We refer to this equilibrium communication strategy simply by "pooling strategy".

[20] This is so as in our model technologies have a common value that is learned before agents communicate in $t = 2$.

[21] In what follows, we assume that the out-of-equilibrium belief $\hat{\pi}_1(Y, Y, Z, Y)$ equals $\hat{\pi}_1(Y, Y, Z, Z)$.

truthful revelation if initial technologies are the same;

(ii) the belief function equals the density $f_1\left(z|\mathbf{I}_1^{\mathsf{dl}}\right) = f(z)$ *for all* z *and* m_2 *in case* $X_{2,1} = Z$*;*

(iii) the decision strategy is a double-cut-off strategy. The cut-off value in case initial technologies are the same, $\bar{y}_{\mathsf{S}}^{\mathsf{gl}}$*, satisfies*

$$\lambda\left[\hat{\pi}_1\left(Y,Y,Y,Y;\bar{y}_{\mathsf{S}}^{\mathsf{gl}}\right) - \hat{\pi}_1(Y,Y,Z,Y)\right] = E\left[\widetilde{Z}|s^Y,s^Y,\bar{y}_{\mathsf{S}}^{\mathsf{gl}}\right] - \bar{y}_{\mathsf{S}}^{\mathsf{gl}}, \tag{7}$$

with $\bar{y}_{\mathsf{S}}^{\mathsf{gl}} \in \left(0,\bar{y}_{\mathsf{S}}^{FB}\right)$ *for* $\lambda < \underline{\lambda}_{\mathsf{dl}}^{\mathsf{gl}}$*.* $\bar{y}_{\mathsf{S}}^{\mathsf{gl}}$ *is a decreasing function of* λ*.*[22] *For* $\lambda \geq \underline{\lambda}_{\mathsf{dl}}^{\mathsf{gl}}$*,* $\bar{y}_{\mathsf{S}}^{\mathsf{gl}} = 0$*. The cut-off value in case initial technologies differ,* $\bar{y}_{\mathsf{D}}^{\mathsf{gl}}$*, satisfies*

$$\lambda\frac{\pi}{1+\pi} = E\left[\widetilde{Z}\right] - \bar{y}_{\mathsf{D}}^{\mathsf{gl}}, \tag{8}$$

with $\bar{y}_{\mathsf{D}}^{\mathsf{gl}} \in \left(0,\bar{y}_{\mathsf{D}}^{FB}\right)$ *for* $\lambda < \bar{\lambda}_{\mathsf{dl}}^{\mathsf{gl}}$*.* $\bar{y}_{\mathsf{D}}^{\mathsf{gl}}$ *is a decreasing function of* λ*. For* $\lambda \geq \bar{\lambda}_{\mathsf{dl}}^{\mathsf{gl}}$*,* $\bar{y}_{\mathsf{D}}^{\mathsf{gl}} = 0$*.*

Figure 3, panels (a) and (b) correspond to (7) and (8), respectively.

Figure 3: Decentralized decision making and global markets. Panels (a) and (b) depict the structure of equilibrium. Panel (c) reports equilibrium cut-off values and reputational gaps for $f_X = 1$ and $\pi = 1/2$. $\Delta\hat{\pi}_1(\mathsf{S})$ denotes $\hat{\pi}_1\left(Y,Y,Y,Y;\bar{y}_S^{\mathsf{gl}}\right) - \hat{\pi}_1(Y,Y,Z,Y)$, and $\Delta\hat{\pi}_1(\mathsf{D}) = \frac{\pi}{1+\pi}$.

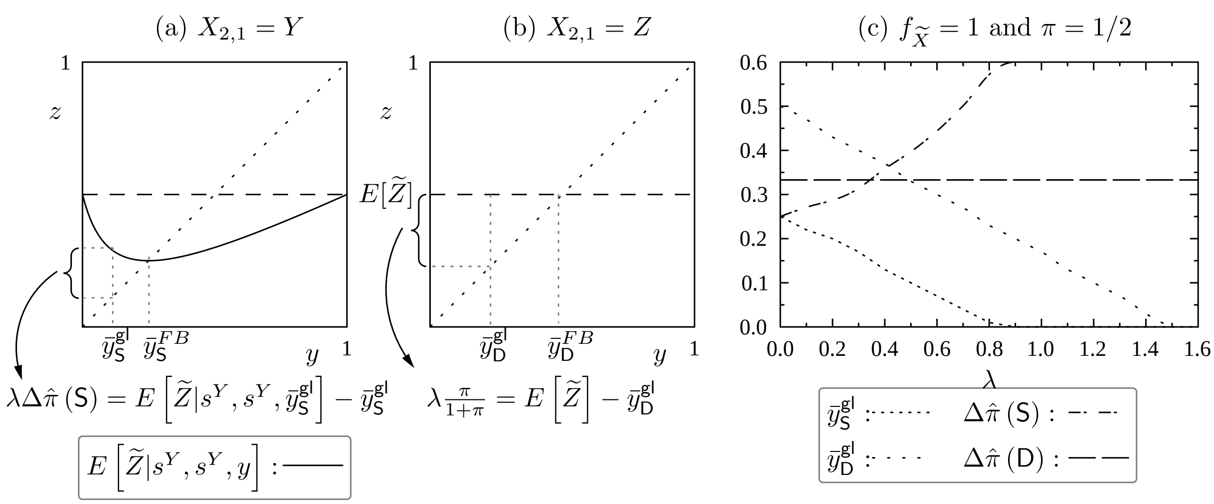

Panel (c) shows the equilibrium values in case of $f_X = 1$ and $\pi = 1/2$. Eq (8) shows that if agents adopted different technologies in $t = 1$, then the reputational gap is a constant

[22] The remark made in footnote 17 concerning $\bar{y}_{\mathsf{ia}}$ applies here to $\bar{y}_{\mathsf{S}}^{\mathsf{gl}}$.

function of $\bar{y}_{\mathsf{D}}^{\mathsf{gl}}$. To understand why, recall that ability means the ability to identify the better technology. When the market observes that agents initially used different technologies, the agents' choices in $t = 2$ either allow the market to infer who used the better and the worse technology (i.e., (Y, Z, Z, Z) and (Y, Z, Y, Y)) or does not allow the market to infer any information on the relative performance of the technologies (i.e., (Y, Z, Y, Z) and (Y, Z, Z, Y)). The value of $\bar{y}_{\mathsf{D}}^{\mathsf{gl}}$ does not provide additional information on an agent's ability. If instead the market observes that agents initially adopted the *same* technology, $\hat{\pi}_1 \left(Y, Y, Y, Y; \bar{y}_{\mathsf{S}}^{\mathsf{gl}} \right)$ does depend on the cut-off value: the lower is $\bar{y}_{\mathsf{S}}^{\mathsf{gl}}$, the lower is the reputation an agent commands in case of continuation.

5 Centralized decision making

Now we turn to centralized decision making. Recall that we assume that $X_{1,1} = Y$. Under centralized decision-making, each agent first reports the value of his technology to the centre. Next, the centre chooses the technology that will be adopted at both sites in $t = 2$. The centre is assumed to maximize $x_{1,2} + x_{2,2}$. Hence, it is an optimal response of the centre to pick the technology with the higher expected value, given the information (I_C) it possesses. More formally,

$$d_C(I_C) = \begin{cases} Y, Y & \text{if } E\left[\widetilde{Y}|\mathbf{I}_C\right] > E\left[\widetilde{Z}|\mathbf{I}_C\right] \\ Y, Y & \text{if } E\left[\widetilde{Y}|\mathbf{I}_C\right] = E\left[\widetilde{Z}|\mathbf{I}_C\right] \text{ and } \mathsf{coin} = Y \\ Z, Z & \text{otherwise} \end{cases} \tag{9}$$

where "coin$= Y$" means that the centre flips a fair coin with faces Y and Z, and Y comes up.[23]

Of course, given this decision strategy, an agent who cares about his reputation does not report the truth, but has an incentive to exaggerate the value of his technology. Exaggeration

[23] We focus on a symmetric treatment of the agents. Asymmetric equilibria are beyond the scope of this paper. Also, one can think of alternative assumptions concerning the decision of the centre in case of $X_{1,1} = Y$ and $X_{2,1} = Z$ and $E\left[\widetilde{Y}|I_C\right] = E\left[\widetilde{Z}|I_C\right]$. For example, the centre could decide that both agents must continue with their initial choices. We don't explore the implications of this alternative assumption in the current paper.

is, however, not without its costs. As a result of it, the centre may decide to impose an inferior technology on both agents. This cost stops an agent from exaggerating too much. As a result, if agents started out with different technologies in equilibrium communication strategies are partition strategies. Such strategies were first described by Crawford and Sobel (1982). In a partition strategy, information is lost as the agent adds noise to his message: he partitions the space of possible technology values $[0,1]$ into $N \geq 1$ intervals, $0 = a_0(N) < a_1(N) < \cdots < a_N(N) = 1$, and reports only to which interval the value of his technology belongs. That is, he ranks its value, and the number of intervals equals the number of possible ranks. We denote a partition by $a(N)$ or simply by a.

In terms of informativeness, a partition strategy is in between the truthful revelation that characterizes communication in case of decentralized decision making with local markets and the absence of communication in case of decentralized decision making process and global markets. That is, the concentration of decision rights deteriorates communication in the former case but improves in the latter. With local reputations, the loss of an agent's decision-making power and its uploading to the centre means that an agent starts to use his communication to indirectly influence the perception of his ability. The quality of communication drops. In case of global markets, the loss of decision-making power makes that an agent becomes cautious when communicating. His exaggerated claims are no longer costless but can lead to an inferior choice at his own site.

Does an agent truthfully report the value of his technology to the centre if the other agent uses the *same* technology in $t = 1$? Agent i's interests are different from those of the centre, but identical to those of agent j. This offers room for the agents to (tacitly) collude, and to induce the centre to choose the technology they deem best. Each can send either of two messages, one, for $y < \bar{y}_{\mathsf{S}}$, such that the centre will next decide that the technology is sufficiently good to merit continuation, and one, for $y > \bar{y}_{\mathsf{S}}$, inducing the centre to force the agents to switch. Note that collusive behaviour of this sort seems easy to sustain as there is no asymmetric information among the agents. Although this is a partition strategy with $N \leq 2$, to distinguish it from the more general partition strategy in case agents use different technologies, we refer to it as a *collusion strategy*. It is completely characterized by a single value, $\bar{y}_{\mathsf{S,cl}} \in [0,1]$, for which an agent is indifferent between sending one message rather than the other.[24] The next proposition sums up. The indifference equations that characterize the partition strategy are stated in the proof.

Proposition 4 *Define* $\bar{\lambda}^{\mathsf{lo}}_{\mathsf{cl}} = E\left[\widetilde{Z}\right]\frac{\left(3+\pi^2\right)(1+\pi)}{4\pi^2}$ *and* $\bar{\lambda}^{\mathsf{gl}}_{\mathsf{cl}} = E\left[\widetilde{Z}\right]\frac{1+\pi^2}{\pi(1+\pi)}$. *In case of centralized decision making, in equilibrium*

(i) the centre's decision strategy is as defined in (9).

(ii) the communication strategy is (a) a partition strategy $(N^*, \mathbf{a}^*)$ *if initial technologies differ*[25]*, and (b) a collusion strategy characterized by* $\bar{y}_{\mathsf{S,cl}}$ *if initial technologies are the same;*

(ii-a) in case of local markets, the partition strategy satisfies $N^* \geq 2$ *for* $\lambda < \bar{\lambda}^{\mathsf{lo}}_{\mathsf{cl}}$*, and* $N^* = 1$ *otherwise; the collusion strategy is characterized by* $\bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}} > 0$ *for* $\lambda < \bar{\lambda}^{\mathsf{lo}}_{\mathsf{cl}}$*, and* $\bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}} = 0$ *otherwise. That is, agents do not send influential information*[26]*, neither on* y *nor on* z*, for* $\lambda \geq \bar{\lambda}^{\mathsf{lo}}_{\mathsf{cl}}$.

(ii-b) in case of global markets, the partition strategy satisfies $N^* \geq 2$ *for any finite* λ*; the collusion strategy is characterized by* $\bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}} > 0$ *for* $\lambda < \bar{\lambda}^{\mathsf{gl}}_{\mathsf{cl}}$*, and* $\bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}} = 0$ *otherwise. That is, in case agents initially used different strategies, agents send influential information about* y *and* z *for any finite* λ*. If agents initially used the same technology, they do not send influential information about the technology's values for* $\lambda \geq \bar{\lambda}^{\mathsf{gl}}_{\mathsf{cl}}$.

The striking result in case of global markets that influential communication is part of equilibrium behaviour for any finite λ stems from the fact that the reputational gap becomes vanishingly small the more the communication strategy of the agent approaches babbling. To see this, consider a partition with $N = 2$. This partition is fully determined by $a_1(2) = a_1$.

[24] Note that had we assumed a collusion strategy under decentralized decision making, there would not have been a difference with the analysis in section 4. In case of local markets, there is no need to collude. In case of global markets, the interests of both agents are perfectly aligned when they start with the same technologies. As a result, they would have taken the same decision concerning period 2 technologies had they been able to collude.

[25] For the exact expressions, see (A3) and (A5).

[26] We say that agent 1 sends *influential information* (or that communication is influential) if there are two messages m_1 and m_1' about Y and a message m_2 about Z such that $d_C(m_1, m_2) = Y$ with probability one and $d_C(m_1', m_2) = Z$ with probability smaller than one. That is, the partition contains at least two intervals, $N \geq 2$.

The equation that determines the value of a_1 becomes[27]

$$\lambda \frac{4\pi}{1+\pi} F(a_1)(1-F(a_1)) = E\left[\widetilde{Z}\right] - a_1. \tag{10}$$

This equation is shown in Figure 4.

Figure 4: Influential communication (i.e., $a_1^* > 0$) for any finite λ in case of centralized decision making and global markets if agents start out with different technologies.

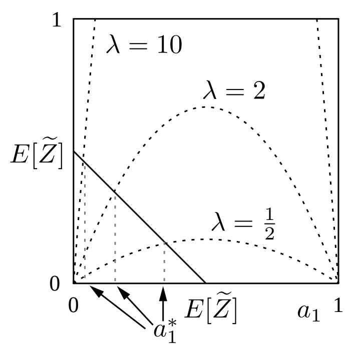

The dotted lines represent the left-hand side of Eq (10) for various values of λ. The drawn line graphs the right-hand side. Clearly, for any finite λ, $a_1^* > 0$ such that influential communication is part of equilibrium behaviour.

To understand why for $a_1 = 0$ the reputational gap $\hat{\pi}_1(Y,Z,Y,Y) - \hat{\pi}_1(Y,Z,Z,Z)$ equals zero, recall that if agents babble about different technologies, the centre decides on the technology to be used by tossing a coin. As a result, the (random) decision of the centre does not add any information on the relative values of the technologies nor on the ability of the agents. Hence, $\hat{\pi}_1(Y,Z,Y,Y) = \hat{\pi}_1(Y,Z,Z,Z)$, and the reputational gap vanishes.

With *local* markets, influential communication about the value of a technology is not possible for $\lambda \geq \bar{\lambda}_{\mathsf{cl}}^{\mathsf{lo}}$ as even for $a_1 = 0$ and $\bar{y}_{\mathsf{S,cl}}^{\mathsf{lo}} = 0$ the reputational gap does *not* vanish. The reason is that it is not known whether agents initially used the same technologies or different ones. If an agent is forced to change technology, it is inferred that agents must initially have used different technologies and that next the centre tossed a coin. The deduced difference in initial technology adoption hurts an agent's reputation. If instead an agent must continue his initial technology this may also mean that both agents initially used the same

[27] See the proof of Proposition 4.

technology. The latter makes it more likely that the agents received a correct signal. As a result, continuation boosts an agent's reputation, and the reputational gap continues to exist even for $a_1 = 0$.

A well-known feature of many cheap-talk models is the existence of multiple equilibria. For instance, in our model a babbling equilibrium in which all messages are ignored by the centre exists for all parameter values. In such an equilibrium, $E\left[\widetilde{Y}|\mathbf{I}_C^{\mathsf{cl}}\right] = E\left[\widetilde{Y}\right]$ and $E\left[\widetilde{Z}|\mathbf{I}_C^{\mathsf{cl}}\right] = E\left[\widetilde{Z}\right]$. A babbling equilibrium, however, is not neologism-proof (Farrel, 1993) in case of globally determined reputations. To see this, suppose $y < E\left[\widetilde{Z}\right]$. Then, agent 1 has an incentive to make clear to the centre that it should not choose technology Y in $t = 2$. As in the babbling equilibrium the centre's decision does not influence the agents' reputations, any neologism meaning "$y < E\left[\widetilde{Z}\right]$" is credible: agent 1 wants to say it if it is true, and does not want to say if it is not true. Neologism-proofness does not completely solve the problem of multiple equilibria. For λ being sufficiently small, equilibria can be distinguished on the basis of the number of partitions. For any partition equilibrium with $N^* > 1$ the reputational gap is positive. As a result, no neologism is credible. That is, neologism-proofness excludes the babbling equilibrium but not equilibria in which meaningful communication takes place. An appeal to Pareto dominance would generally lead to the selection of the equilibrium with the largest number of partitions.[28]

6 Welfare Comparisons

We now turn to the main question of the paper: how is the quality of the learning process affected by the assignment of decision rights and the information on which markets base the reputations of the agents? To answer this question, we fix a decision process p, fix the way reputations are determined, fix the values of λ and π, and assume $s_1 = s^Y$. The previous sections then determine equilibrium behaviour. Before 1 observes y, and given equilibrium behaviour what is the expected value of the technology in use at site 1 in $t = 2$? We denote

[28] Chen et al. (2008) have proposed NITS – no incentive to separate – as an equilibrium refinement: "An equilibrium satisfies NITS if the Sender of the lowest type weakly prefers the equilibrium outcome to credibly revealing his type (if he somehow could)" (p. 118). The babbling equilibrium does not satisfy NITS.

this *ex ante* expected *ex post* value by $E\left[\widetilde{X}_{1,2}|s^Y,\lambda,\pi\right]$. The theoretical maximum value is $E\left[\widetilde{Y}|y>z\right]$, which obtains if agent 1 chooses the better technology in period 2 with probability one. No process generates this value, unless $\pi=1$ in which case the better technology is identified in $t=1$. The "theoretical minimum" value is $\pi E\left[\widetilde{Y}|y>z\right]+(1-\pi)E\left[\widetilde{Y}\right]$. This is the expected value in case the technology adopted at site 1 in $t=2$ equals the first period choice with probability one, independent of the experience gained with the technologies in $t=1$ throughout the economy.

To focus on differences in value creation thanks to learning from own past behaviour and from the experience of others, we transform $E\left[\widetilde{X}_{1,2}|s^Y,\lambda,\pi\right]$ using the following formula,

$$W(\lambda,\pi)=\frac{E\left[\widetilde{X}_{1,2}|s^Y,\lambda,\pi\right]-\left(\pi E\left[\widetilde{Y}|y>z\right]+(1-\pi)E\left[\widetilde{Y}\right]\right)}{E\left[\widetilde{Y}|y>z\right]-\left(\pi E\left[\widetilde{Y}|y>z\right]+(1-\pi)E\left[\widetilde{Y}\right]\right)}*100\%. \qquad (11)$$

That is, $W(\lambda,\pi)\in[0\%,100\%]$ captures value creation thanks to learning, over and above the minimum value, as a percentage of what is maximally attainable. We refer to it as 'welfare.'

6.1 Decentralized decision making: welfare comparisons

In this subsection and the next, we analyse how welfare changes if, for a given assignment of decision rights, the information that markets have changes from local to global. The main finding is that outcomes are generally better when markets are local. However, reputational concerns may distort decision making under global markets so severely that sometimes global markets perform better. Thus, more information in the hands of reputation-driven agents may give rise to a reduction in welfare.

Key to welfare comparisons are (i) the information agents have, and (ii) the degree to which they use it in the various situations. Consider (i). By definition, an isolated agent only knows the value of his own technology, and does not know what technology has been adopted at the other site. We know from Propositions 2 and 3 that with global markets agent 1 also knows the technology $X_{2,1}$ used by the other agent (but not its value $x_{2,1}$ if $X_{2,1}=Z$), and that if markets are local he knows both the technology and its value, $X_{2,1}$ and $x_{2,1}$.

Consider (ii). Obviously, if an agent does not care about his reputation, any additional information can only lead to an increase in welfare. Hence, for low values of λ, $W_{\mathsf{ia}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{gl}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda,\pi)$, and outcomes are best when markets are local.[29] However, with strong reputational concerns more information is not always better in case agents can learn from the experience of others. The propositions establish the values of λ above which an agent ignores all information and simply continues with his initial choice of technology. These values are $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}} = 1/\pi$ and $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{gl}} = E\left[\widetilde{Z}\right](1+\pi)/\pi$ for an agent whose reputation is locally or globally determined, respectively. For $E\left[\widetilde{Z}\right](1+\pi) > 1$, $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{gl}} > \bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$ holds.[30] In other words, if this condition is satisfied, for sufficiently high values of λ, welfare goes up when markets are global, although agents have communicated less information among each other.[31]

To better understand the difference in welfare in the two situations, we first isolate the effect of the information that *markets* have about agents by assuming that agents do not share information. Then we study the effect of additional information that *agents* have in equilibrium in case of local markets.[32] We depart from the situation in which agents initially started with different technologies. We focus on the threshold values of λ above which agents simply continue with their initial technologies. Suppose that there is no communication. Then, agent 1 compares his observed value y with $E\left[\widetilde{Z}\right]$. The reputational gap is narrower under global markets because global markets observe that agents adopted different technologies.[33] This damages the agents' reputations. Thus, the value of λ above which an agent continues with his initial technology is higher than in case of local markets. It is in this sense that, keeping the communication strategy fixed, the incentive to distort the decision on the project is smaller in case of global markets. There is however a second difference

[29] Recall that ia stands for isolated agents, dl for decentralized decision-making, and gl (lo) for globally (locally) determined reputations.

[30] Note that the 1 in $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}} = 1/\pi$ is the upperbound of the support of $f(\cdot)$. The inequality therefore holds independent of the chosen support.

[31] As $\bar{\lambda}_{\mathsf{ia}} = E\left[\widetilde{Z}\right]/\pi < \bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}, \bar{\lambda}_{\mathsf{dl}}^{\mathsf{gl}}$ for all parameter values, an isolated agent stops using information for a lower value of λ than an agent who can learn from another agent.

[32] We thank an anonymous referee for suggesting this approach.

[33] In case of locally determined reputations, $\hat{\pi}_1(Y,Y) = \pi$, while $\hat{\pi}_1(Y,Z) = 0$ is a plausible out-of-equilibrium belief. If instead markets are global, then it follows from (8) that the reputational gap equals $\frac{\pi}{1+\pi}$.

between local and global markets: in equilibrium in local markets the agents can truthfully share their private information with each other. As a result, the difference $|z - y|$ can be as large as one, meaning that the value of λ above which an agent is willing to continue with his initial technology even if $y = 0$ goes up and becomes $\lambda = 1/\pi$. Clearly, if $E\left[\widetilde{Z}\right](1+\pi) > 1$ the net effect of a larger reputational gap and more information exchange in equilibrium is such that for high values of λ, local markets create less welfare than global markets. This inequality holds if it is sufficiently easy to identify the better technology (π high), and if the unconditional expected value $E\left[\widetilde{Z}\right]$ of a technology is sufficiently high. In case of the uniform distribution or any other symmetric distribution it cannot hold.[34]

In Figure 5, we depict W, for decentralized decision making with local or global markets and for isolated agents under the assumption that the value of technology $X \in \{Y, Z\}$ is uniformly distributed, $f_X(x) = 1$ on $[0, 1]$, $\pi = \frac{1}{2}$, and $\lambda > 0$.[35]

Figure 5: $W(\lambda, \pi)$ for isolated agents and for decentralized learning from others with local and global markets. $f_{\widetilde{X}} = 1$ and $\pi = 1/2$ such that $\bar{\lambda}_{\mathsf{ia}} = 1$, $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{gl}} = 3/2$, and $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}} = 2$.

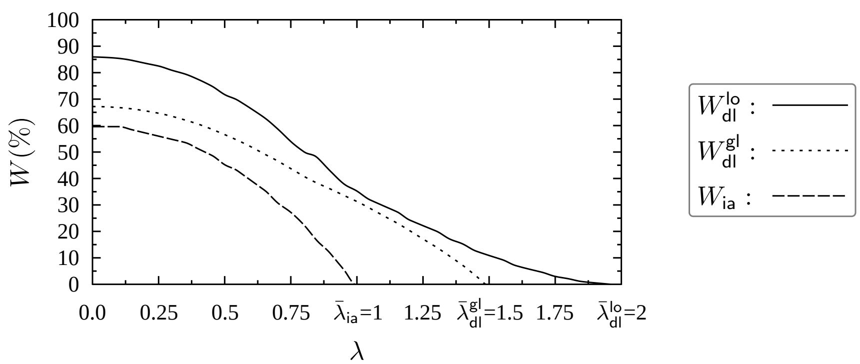

Figure 5 illustrates a number of points. First, for low values of λ, local markets generate

[34] Interestingly, a similar sort of less-is-more effect involving information transmission is present in Fisher and Stocken (2001) and Ivanov (2010). They establish that in a single-Sender, single-Receiver game less information in the hands of the Sender may improve communication with the Receiver.

[35] For $f_X = 1$ and $\pi = 1/2$, $E\left[\widetilde{Y}|y > z\right] = 2/3$ and $\pi E\left[\widetilde{Y}|y > z\right] + (1-\pi) E\left[\widetilde{Y}\right] = 7/12$. Note that $W_{\mathsf{ia}}(\lambda, \pi)$ and $W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda, \pi)$ are (right) continuous in λ at $\lambda = 0$ as equilibrim behaviour changes continuously for λ. In case of global markets, however, communication behaviour changes discontinuously (from truth telling for $\lambda = 0$ to no communication for $\lambda > 0$ and $X_{2,1} = Z$). As a result, $W_{\mathsf{dl}}^{\mathsf{lo}}$ drops discontinuously at $\lambda = 0$. Furthermore, $W_{\mathsf{dl}}^{\mathsf{lo}}(0, \pi) = W_{\mathsf{dl}}^{\mathsf{gl}}(0, \pi)$ for all π. Finally, $W_{\mathsf{dl}}^{\mathsf{lo}}(0, \pi) = W_{\mathsf{dl}}^{\mathsf{gl}}(0, \pi) < 1$ as $\pi < 1$.

more welfare than global markets. Second, learning from one's own past behaviour and from others potentially boosts welfare enormously. For λ close to zero, and $\pi = 1/2$, an isolated agent can capture 60% of the increase in expected project value. Learning from others further increases this percentage. Third, if *markets* can compare agents' behaviour across sites, this reduces the positive effect of learning from others. The main reason is that communication breaks down when markets can make comparisons.[36]

Proposition 5 *For any f_X and π, there exists a $\lambda_1 > 0$ such that $W_{\mathsf{ia}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{gl}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda,\pi)$ for all $\lambda < \lambda_1$. Furthermore, for any f_X and π such that $E\left[\widetilde{Z}\right](1+\pi) < 1$, there exists a $\lambda_2 > 0$ such that $W_{\mathsf{ia}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{gl}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda,\pi)$ for all $\lambda > \lambda_2$. If instead f_X and π satisfy $1 < E\left[\widetilde{Z}\right](1+\pi)$, then there exists a $\lambda_3 > 0$ such that $W_{\mathsf{ia}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{gl}}(\lambda,\pi)$ for $\lambda > \lambda_3$. For $f_X = 1$, the uniform distribution, $W_{\mathsf{ia}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{gl}}(\lambda,\pi) < W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda,\pi)$ holds for all λ and π.*

6.2 Centralized decision making: welfare comparisons

The key finding in this section is that for high values of λ welfare is higher if the centre receives her information from agents whose reputations are *globally* determined. This result stems from the fact, reported in Proposition 4, that with global markets, agents send influential information for all values of the parameters, while they will babble for sufficiently high values of λ if markets are local.

Proposition 6 *For any f_X and π there exists a $\lambda_4 > 0$ such that $W_{\mathsf{ia}}(\lambda,\pi) < W_{\mathsf{cl}}^{\mathsf{lo}}(\lambda,\pi) < W_{\mathsf{cl}}^{\mathsf{gl}}(\lambda,\pi)$ for all $\lambda > \lambda_4$.*

Another finding is that if agents used the same technology, information transmission takes place for a wider range of parameters under local markets, $\bar{\lambda}_{\mathsf{cl}}^{\mathsf{lo}} > \bar{\lambda}_{\mathsf{cl}}^{\mathsf{gl}}$. The reason for this result is that global markets observe that agents intially used the same technology. This boosts the agents' reputations. When more reputation is at stake, agents have weaker incentives to change technology.

[36] If $\lambda \in \left(\bar{\lambda}_{\mathsf{dl}}^{\mathsf{gl}}, \bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}\right)$, then $W_{\mathsf{ia}}(\lambda,\pi) = W_{\mathsf{dl}}^{\mathsf{gl}}(\lambda,\pi) = 0 < W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda,\pi)$. If $\lambda \geq \bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$, then, $W_{\mathsf{ia}}(\lambda,\pi) = W_{\mathsf{dl}}^{\mathsf{gl}}(\lambda,\pi) = W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda,\pi) = 0$. These cases are ignored in Proposition 5.

Figure 6: $W(\lambda, \pi)$ for isolated agents and for centralized learning from others with locally and globally determined reputations. $W(\lambda, \pi)$ in case of centralized learning is based on a partition strategy with at most two ranks. $f_{\tilde{X}} = 1$ and $\pi = 1/2$, such that $\bar{\lambda}_{\mathsf{ia}} = 1$, $\bar{\lambda}_{\mathsf{cl}}^{\mathsf{lo}} = 2\frac{7}{16}$.

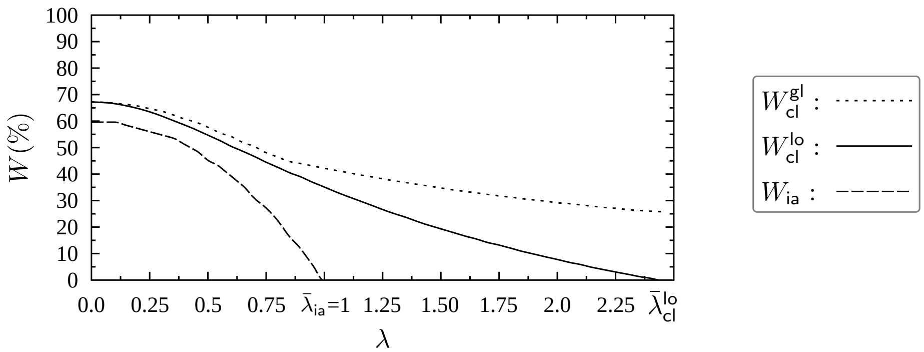

Besides, we established numerically that for the uniform distribution, $W_{\mathsf{ia}}(\lambda, \pi) < W_{\mathsf{cl}}^{\mathsf{gl}}(\lambda, \pi), W_{\mathsf{cl}}^{\mathsf{lo}}(\lambda, \pi)$ holds for all $\lambda > 0$ and π. Proposition 6 is illustrated in Figure 6 for the uniform distribution and $\pi = 1/2$. To simplify calculations, we have imposed that communication with the centre is limited to at most two ranks in case agents initially used different technologies. Clearly, if agents can learn from others welfare improves. Because of our limitation to at most two ranks, the graph understates the benefits for low values of λ. In fact, for $\lambda = 0$, agents would truthfully reveal their private information and the performance of a centralized learning process would equal that of a decentralized learning process. We then know from Figure 5 that $W \approx 86\%$ rather than $W \approx 68\%$ as shown in the graph. The graph nicely illustrates the key finding for high values of λ: welfare is higher when markets observe technology adoption at both sites, as communication between agents and centre remains influential for any finite λ, whereas it dies out for high values of λ in case of locally determined reputations.

6.3 A final comparison

In the previous two subsections we have analysed how welfare changes if, for a given assignment of decision rights, the information on which the perceptions of agents' abilities are based changes from local to global. In this subsection we present a complementary result: with global markets, welfare is higher under centralized than under decentralized decision making.

Proposition 7 *In case of global markets, and for any f_X, π, and λ, welfare $W(\lambda, \pi)$ is higher with centralized than with decentralized decision making.*

The main benefit of moving from a decentralized process to a centralized one in case of global markets is the restoration of communication when agents initially used different technologies. The proof establishes that even if agents in a centralized decision process were to limit themselves to a communication strategy consisting of at most two ranks - and choose the partition optimally - welfare goes up. This suggests that the welfare difference can be substantial for low values of λ, as such values allow for richer communication (i.e., finer partitions).[37]

If under global markets agents used the same technology, a decentralized process and a centralized process yield equal outcomes. The threshold strategy under a decentralized process coincides with the collusion strategy under a centralized process. When the markets know that the agents used the same technologies, reputation formation does not depend on the type of decision-making process. A local market, on the other hand, does not know which technology is used by the other agent. As a result, reputation formation does depend on the decision-making process, and the threshold strategy under a decentralized process does not coincide with the collusion strategy under a centralized process.

7 Discussion

7.1 Two examples of the struggle with learning from others

An important objective of this paper was to gain insight into the effects of alternative learning processes on the quality of decisions in situations where information is dispersed among agents, and agents are concerned about their reputations. Our analysis focuses on two broad features of decision-making processes: whether a decision-making process is centralized or decentralized, and whether reputations of decision-makers are based on local information only, or can also be based on comparisons across sites. We believe that our model describes

[37] If reputations are locally determined, the learning process that is best depends fundamentally on the parameters of the model, see an earlier working paper version of this paper, Swank and Visser 2010.

a number of important dimensions and therefore sheds light on variety of real-world learning processes.

For example, consider the medical profession. The delivery of medical interventions varies widely from place to place.[38] This variation has been a source of worries as, most likely, some patients do not receive optimal treatment.[39] It also offers scope for learning. In response, physicians' associations and health care authorities have exerted much effort to design learning processes in which locally gained experiences are compared, and best practices – interventions, surgical procedures, drug use – diffused. In the medical sector, expert panels are frequently used to evaluate the evidence on the effectiveness of rival practices in a given field. Given the close ties between experts and industry, and the long gestation period that characterizes the development of practices, experts tend to identify with certain practices. The result, according to students of expert panels, is "process loss" due to reputational concerns, leading in turn to poor information exchange and aggregation in the meetings, and a low adoption rate of best practices afterwards.[40] Organizing these panels is therefore fraught with problems. An important organizational dimension is the degree of freedom individual physicians have in following the outcomes of panel meetings - the decision rights dimension in our model. Eddy (1990) distinguishes, in increasing degree of freedom, standards, guidelines, and options.[41] It also seems that the IT revolution and increased information dissemination over the internet, in combination with societal pressure to increase disclosure makes it easier for patients and authorities to compare medical practices across places. This information can then shape the perception of physicians' abilities - the local versus global reputations.

In terms of our model, if the environment becomes more global and allows for comparisons across physicians or hospitals, learning processes that are decentrally organized suffer from poor information exchange and low adoption rates of best practices. Centralized processes fare better. Expert panels should be empowered with the authority to impose standards.

The European Union is another case in point. It has been promoting the so-called open method of coordination (OMC) to foster learning and the diffusion of best practices in many

[38] That variation is large is a well-established fact, see Phelps (2000).

[39] See, e.g., Eddy (1990).

[40] See Fink et al. (1984) and Rowe, Wright, and Bolger (1991).

[41] Eddy (1990) distinguishes, in increasing degree of freedom, standards, guidelines, and options.

policy areas. The hope is that goals like EU competitiveness can be furthered by avoiding the grand questions about the best model for Europe and by taking instead a more pragmatic micro-orientation in which countries that face similar problems seek to learn from each other.[42] Rather than relying on legislation by Brussels–a form of centralized decision-making–, the OMC leaves decision rights with the EU countries: they decide whether to implement the lessons learned. Moreover, instead of applying formal sanctions to transgressors, the OMC turns to naming and shaming to expose a country's weak performance in public, and applies peer pressure if a country opposes adoption of superior policies.[43] In practice, the method is not considered to be very successful in guaranteeing a high quality learning process. It is generally felt that countries exaggerate the success of their current practices. Also, the implementation of new ideas is very limited. Claudio Radaelli (2003, p. 12) argues that these disappointing results stem from a misguided view of policy makers among the proponents of the OMC. Rather than caring about the truth, they care about 'political capital' and 'prestige' – forms of reputational concerns. Arguably, the 'naming and shaming in public' suggests that the perception of an agent's ability in the case of the OMC can be based on comparisons across countries.

The OMC is a decentralized decision-making process in an environment where reputations are based on global information. Our model shows that in such an environment one may expect that countries exaggerate the success of their current practices and are reluctant to adopt new practices. In a global world, centralization may facilitate communication and the adoption of best practices.

7.2 First period behaviour

Our analysis has focused on what agents report about their locally gained experiences at the end of period one and on decision making in period two. We now briefly revisit the imputed strategy of following one's signal in period one. Assume that the communication strategies and second period behaviour are as described in the previous sections. Suppose

[42] The OMC has been applied in areas as diverse as employment, social inclusion, innovation, education, occupational health and safety.

[43] See Pochet (2005) and Radaelli (2003).

moreover that agent 2 follows his signal. Clearly, if agent 1 were to deviate from his strategy of following his signal, he would reduce the expected value of the technology adopted in period one. Moreover, by following his signal, an agent minimizes the probability that he wants to change or is forced to change his technology in period 2, and thus the probability that his reputation is hurt. By following his signal, he maximizes the probability that the technology he adopts in period 1 is the same as the one adopted by agent 2. In case of global markets, this maximizes his expected reputation. Finally, by following his signal he provides the best input to the centre on which she bases the technology choice in period 2, thus maximizing expected second period payoff stemming from the adopted technology. In short, following his signal is part of equilibrium.

The strategy of following one's signal is the root cause that a reputational gap exists, and therefore that in the continuation of the game reputational concerns get in the way of taking the decision that maximizes the expected value of the technologies adopted. Had an agent ignored his signal and made a random choice in period one, no information about his ability could have been deduced from a comparison of observed technologies across agents or from a sequence of choices. However, an equilibrium in which agents ignore their signals and instead choose randomly in $t = 1$ does not exist: in the absence of a reputational gap, an agent wants to deviate by following his signal, as this increases the expected payoff in the first and second period without hurting his reputation.

An equilibrium may exist in which regardless of signals agent 1 adopts Y and agent 2 adopts Z in period 1. In such an equilibrium, neither agent has incentives to report dishonestly or to distort decisions in period 2 (see below). However, agent 1 may have an incentive to deviate in period 1 if he received a signal indicating that technology Z is the better technology. Then, switching to technology Z in period 1 would increase period 1 payoff. The cost of switching is that markets are likely to infer from switching that agent 1 has received $s_1 = s^Z$. A direct implication is that agent 1's reputation will be at stake in period 2. Thus, in period 1, agent 1 who has received $s_1 = s^Z$ faces a trade-off. From a period 1 perspective he should adopt technology Z, but from a period 2 perspective he should keep technology Y. Obviously, adopting Z is attractive if π is high, λ is low, and period 1 is important relative to period 2.

7.3 Avoiding distortive behaviour in period 2

One way of avoiding distorting behaviour in period 2 is fixing technologies before the agents receive their signals in period 1. For instance, the agents can decide that in period 1 agent 1 adopts technology Y and agent 2 adopts technology Z. Alternatively, a centre may allocate technologies in this way in period 1. This division of technologies allows for perfect learning in period 2. Because the allocation of technologies does not depend on signals, reputational concerns do not hinder communication or decision making. As a result, in period 2 both sites always adopt the better technology, yielding a payoff equal to $E[Y|y > z]$. Fixing technologies in period 1 yields a higher payoff in period 2 than not fixing technologies for two reasons. First, reputational concerns do not distort communication or decision making. Second, by allocating technologies in the right manner, agents always learn both y and z. In our model, the agents may only learn the value of y. This diminishes the scope of learning. The cost of fixing technology in period 1 is ignoring possible information in period 1. This cost realizes if both agents receive the same signal. Note that if both agents receive different signals, the allocation of technologies is the same as when technologies were fixed. The cost of fixing technologies equals $2\Pr\left(s^Y, s^Y\right)\left[E\left[Y|s^Y, s^Y\right] - E[Y]\right]$. This cost is an increasing function of π. The higher is the probability that the agents are smart, the higher is the likelihood that they receive the same signal, and that these signals are correct. In the extreme that $\pi = 1$, the cost of fixing technologies is highest. The benefits of fixing technologies fall in period 2. The benefit of guaranteeing communication and proper decision making is increasing in λ. Stronger reputational concerns lead to larger distortions in period 2. The effect of π on the benefit of fixing technologies consists of two parts. First, a higher value of π widens the reputational gap. This increases the benefit of fixing technologies. Second, a higher value of π reduces the need for learning. In the extreme that $\pi = 1$, there is no need for learning in period 2 at all. Learning should occur in period 1. Technologies should not be fixed. Finally, the decision whether or not to fix technologies in period 1 depends on time preferences. As the cost of fixing technologies fall in period 1 and the benefits fall in period 2, fixing technologies is especially attractive if outcomes in period 2 are important relative to outcomes in period 1.

7.4 Alternative assumptions and future research

Although our model may help to identify the conditions under which learning from others benefits from centralized or decentralized decision making, it is based on a number of restrictive assumptions. We conclude this paper by briefly discussing some of these assumptions.

Centralization. One important assumption is that in a centralized process the centre always acts in the general interest. This is a common assumption in the literature that compares centralized and decentralized decision making, see e.g. Alonso et al (2008). In reality, there is little reason to put so much confidence in central bodies or corporate head quarters. For example, a centre may be biased towards one of the technologies because of favoritism. Alternatively, a centre may be biased because somehow its name is connected to one of the technologies. Of course, our assumption of a "benevolent" centre provides too favourable a picture of centralized processes. An analysis of the repercussions of a biased centre on behaviour of both agents, both advantaged and disadvantaged, is an interesting topic for further research.

Information that agents have. We have described the private information that agents have as non-verifiable, and communication as cheap talk. Although this may well reflect an important part of information agents have gained locally, they may also have verifiable information. Such information can be checked by other agents. If it is unknown whether an agent actually possesses information that is decision-relevant to another agent, the former may have an incentive to selectively withhold his private information from the latter, see e.g. Milgrom and Roberts (1986). How does the presence of verifiable information change our findings? Although the nature of information manipulation changes, the incentives to manipulate continue to be determined by the interplay of the decision rights and the information on which reputations are based. As a result, the quality of information exchange depends in essentially the same way on these same two factors. Consider decentralized decision-making with local markets. The fact that an agent's reputation is independent of what the other agent does and that an agent can decide himself what technology he uses next makes that revealing all positive and negative pieces of information is an equilibrium strategy. If reputations are also based on comparisons across sites (and authority remains decentralised), it is important from a reputational perspective to convince the other agent to switch to "your" technology. As a result, any negative information will be withheld. The introduction of centralised decision-making in such a situation gives rise to the selective revelation of negative information. On the one hand, as the agent at a site loses decision-making power, he wants to make sure that the centre is well-informed. On the other hand, his reputational concerns imply that he wants the centre to impose "his" technology at either site. Ceteris paribus, the more damaging negative information is for the technological value, the more likely it is that the information is revealed. Similarly, the more damaging negative information is for his reputation, the less likely it becomes that this information is revealed.

Information that markets have. As in Prendergast and Stole (1996), we assume that markets only observe the adopted technologies. Anecdotal evidence suggests that this captures well certain situations in, e.g., the medical sector. It often suffices to speak with a foreign colleague to realize whether medical practices differ between countries. But it is quite a different matter to find out which of a number of practices is the better one. It is however still an interesting question to ask what would happen if markets do receive an imperfect signal about the performance of a technology before the technology to be adopted in $t = 2$ is chosen. An analysis of such a situation is well beyond the scope of this paper and must await further research.

Appendix A

Proposition A1 *First-best behaviour is not part of equilibrium, whether decision making is centralized or decentralized and whether markets are local or global.*

Proof: We discuss each of the four cases in turn. (i) dl and lo. With first-best behaviour, $\hat{\pi}_1(Y, Y; 0, 0) > \hat{\pi}_1(Y, Z; 0, 0)$. Then, for any $\lambda > 0$, there is a value of ε satisfying $0 < \varepsilon < \lambda(\hat{\pi}_1(Y, Y; 0, 0) - \hat{\pi}_1(Y, Z; 0, 0))$ such that if $y \in (m_2 - \varepsilon, m_2)$ agent 1 benefits from deviating from the first-best decision strategy by sticking to Y.

(ii) dl and gl. Suppose $X_{2,1} = Z$. In case of first-best behaviour, in $t = 2$, either both agents adopt Y or both agents adopt Z. Then, $\hat{\pi}_1(Y, Z, Y, Y) = \frac{2\pi}{1+\pi} > \pi$ and $\hat{\pi}_1(Y, Z, Z, Z) = 0$.[44]

[44] See the proof of Proposition 3 for the derivation. Intuitively, in the former (latter) case the market infers

Note that a unilateral deviation in the decision stage to a pattern of technology adoption the reputation of which is not determined by the imputed equilibrium behaviour. However, it is consistent with the model to assume that any shift to Y in $t=2$ boosts agent 1's reputation, while any shift to Z boosts agent 2's reputation.

Assumption A1 *Consider any adoption vector with $X_{1,1}=Y$. The reputation of* 1 *increases if* 1 *(resp.* 2*) changes from $X_{1,2}=Z$ to $X_{1,2}=Y$ (resp. from $X_{2,2}=Z$ to $X_{2,2}=Y$).*

Then, for every λ there is a value of ε that satisfies $0<\varepsilon<\lambda\hat{\pi}_1(Y,Z,Y,Z)$ such that if $y\in(m_2-\varepsilon,m_2)$ agent 1 benefits from deviating from the first-best decision strategy by sticking to Y.

(iii) cl and lo. The first-best decision strategy of the centre, (9), and truthful reporting by the agents private information imply that $\hat{\pi}_1(Y,Y)>\hat{\pi}_1(Y,Z)$. Thus, for every $\lambda>0$, there is a value of ε that satisfies $0<\varepsilon<\lambda[\hat{\pi}_1(Y,Y)-\hat{\pi}_1(Y,Z)]$ such that agent 1 benefits from deviating from the first-best communication strategy by using $m_1(y)=y+\varepsilon$.

(iv) cl and gl. The same line of reasoning as under (iii) applies.

Appendix B

Proof of lemma 1: Consider (1) in the text. (a) As $\Pr\left(\bar{\theta}|s^Y,0\right)=0$, $E\left[\widetilde{Z}|s^Y,0\right]=E\left[\widetilde{Z}\right]$. Similarly, as $\Pr\left(\bar{\theta}|s^Y,1\right)=1$, then $E\left[\widetilde{Z}|s^Y,1,\bar{\theta}\right]=E\left[\widetilde{Z}\right]$, and therefore $E\left[\widetilde{Z}|s^Y,1\right]=E\left[\widetilde{Z}\right]$. Moreover, $E\left[\widetilde{Z}|s^Y,y,\bar{\theta}\right]<E\left[\widetilde{Z}\right]$ for $y\in(0,1)$, as the term on the LHS is the expected value of the truncated distribution on $[0,y)$. (b) To determine the derivative, use Bayes' rule to write $\Pr\left(\bar{\theta}|s^Y,y\right)=2F(y)\pi/(2F(y)\pi+(1-\pi))$. Also, $E\left[\widetilde{Z}|s^Y,y,\bar{\theta}\right]=\int_0^y tf(t)\,dt/F(y)$. One can verify that $\partial\Pr\left(\bar{\theta}|s^Y,y\right)/\partial y=\Pr\left(\bar{\theta}|s^Y,y\right)\left(1-\Pr\left(\bar{\theta}|s^Y,y\right)\right)\frac{f(y)}{F(y)}>0$, and that $\partial E\left[\widetilde{Z}|s^Y,y,\bar{\theta}\right]/\partial y=\left(y-E\left[\widetilde{Z}|s^Y,y,\bar{\theta}\right]\right)\frac{f(y)}{F(y)}$. Hence,

$$\partial E\left[\widetilde{Z}|s^Y,y\right]/\partial y=\Pr\left(\bar{\theta}|s^Y,y\right)\frac{f(y)}{F(y)}\left(y-E\left[\widetilde{Z}|s^Y,y\right]\right),$$

from which it follows immediately that $E\left[\widetilde{Z}|s^Y,y\right]$ is decreasing for $y<E\left[\widetilde{Z}|s^Y,y\right]$ and increasing for $y>E\left[\widetilde{Z}|s^Y,y\right]$. Hence, $y=E\left[\widetilde{Z}|s^Y,y\right]$ has a unique solution. ■

---

that in $t=1$ agent 1 adopted the better (worse) technology.

Proof of Proposition 1: First, $\hat{\pi}(YY;\bar{y}) = \Pr(\bar{\theta}|YY;\bar{y})$ in (2). Use $\Pr(YY|\bar{\theta}) = \Pr(y \geq \bar{y}|\bar{\theta})/2 = (1-F(\bar{y})^2)/2$ and $\Pr(YY|\underline{\theta}) = \Pr(y \geq \bar{y}|\underline{\theta})/2 = (1-F(\bar{y}))/2$, and apply Bayes rule (analogously for $\hat{\pi}(YZ;\bar{y})$. Clearly, for given reputations the equilibrium strategy is a single threshold strategy with $\bar{y}^{\mathsf{ia}}$ satisfying (3). Given this strategy, equilibrium reputations are as in (2) with $\bar{y} = \bar{y}^{\mathsf{ia}} \leq \bar{y}_{\mathsf{ia}}^{FB}$. To see that $\bar{y}^{\mathsf{ia}}$ is a decreasing function of λ for $\lambda \leq \bar{\lambda}_{\mathsf{ia}}$, define $\delta := \bar{y}_{\mathsf{ia}}^{FB} - \bar{y}_{\mathsf{ia}}$ and $\Delta\hat{\pi} := \hat{\pi}(YY) - \hat{\pi}(YZ)$. Then $(\delta, \Delta\hat{\pi}) \in L := \left[0, E\left[\widetilde{Z}\right]\right] \times [0,1]$, and so L is a complete lattice. It follows from Lemma 1 that (3) can be written as $\delta = k_1(\Delta\hat{\pi}, \lambda)$. It follows from (3) that the function k_1 satisfies $\partial k_1/\partial\Delta\hat{\pi}, \partial k_1/\partial\lambda > 0$, and from (2) that $\Delta\hat{\pi} = k_2(\delta)$ is an increasing function of δ. Hence, we can apply Theorem 3 in Milgrom and Roberts (1994). The set of fixed points of $k : L \times \mathbb{R}^+ \to L$ is non-empty and equals the set of equilibria, and $\delta^* = \bar{y}_{\mathsf{ia}}^{FB} - \bar{y}^{\mathsf{ia}}$ is increasing in λ. Moreover, in case this set is not a singleton, both the highest and the lowest fixed point are increasing in λ. It is straightforward to check that for $\lambda \geq \bar{\lambda}_{\mathsf{ia}}$, $\bar{y}^{\mathsf{ia}} = 0$. It is then plausible to define the out-of-equilibrium $\hat{\pi}(YZ;0) = \lim_{\bar{y}_{\mathsf{ia}} \to 0} \hat{\pi}(YZ;\bar{y}^{\mathsf{ia}}) = 0$. ■

Proof of Proposition 2: The equilibrium belief functions follow immediately from the equilibrium message strategies. That the decision strategy is a double-threshold strategy follows from the analysis preceding the statement of the proposition. Finally, note that for $\bar{y}_{\mathsf{S}}^{\mathsf{lo}} = 0$, the RHS of (5) equals $E\left[\widetilde{Z}\right]$, and therefore $\bar{y}_{\mathsf{D}}^{\mathsf{lo}} = E\left[\widetilde{Z}\right]$, and thus $\lambda = \underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$. Finally, if $\bar{y}_{\mathsf{D}}^{\mathsf{lo}} = 0$ and $\bar{t}_{\mathsf{D}}^{*} = 1$, $\hat{\pi}(YY;0,1) = \pi$ (as agent uses pooling strategy) and $\hat{\pi}(YZ;0,1) = 0$ (this is an out-of-equilibrium belief, the limit of $\hat{\pi}(YZ)$ in case $\bar{y}_{\mathsf{D}}^{\mathsf{lo}} \uparrow 1$) such that for $\lambda \geq \bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$, the agents indeed continue with their initial technologies no matter what.■

The next lemma proves useful in the proof of Proposition 3.

Lemma A1 $E\left[\tilde{Z}|s^Y, s^Z, \bar{y}, z \in A\right] = E\left[\tilde{Z}|z \in A\right]$ *for all* $\bar{y} \in [0,1]$ *and* $A \subseteq [0,1]$.

Proof: Let a given message m_2 be sent for $z \in A$. To prove that $E\left[\tilde{Z}|s^Y, s^Z, \bar{y}, z \in A\right] = E\left[\tilde{Z}|z \in A\right]$ for all $(\bar{y}, A)$, it suffices to show that $f\left(z|s^Y, s^Z, \bar{y}, z \in A\right) = \frac{f(z)}{\int_A f(z)dz} I_A(z)$ for all $(\bar{y}, A)$, where $I(\cdot)$ is the indicator function and $y = \bar{y}$ is the observed value of y. Let $\mathbf{\Omega} := \left\{s^Y, s^Z, \bar{y}, z \in A\right\}$, and write $f(z|\mathbf{\Omega}) = f(z|\mathbf{\Omega}; sm, db)\Pr(sm, db|\mathbf{\Omega}) + f(z|\mathbf{\Omega}; db, sm)\Pr(db, sm|\mathbf{\Omega}) +$

$f(z|\mathbf{\Omega}; db, db)\Pr(db, db|\mathbf{\Omega})$.

$$\begin{aligned} f(z|\mathbf{\Omega}) &= f(z|\mathbf{\Omega}; sm, db)\Pr(sm, db|\mathbf{\Omega}) + f(z|\mathbf{\Omega}; db, sm)\Pr(db, sm|\mathbf{\Omega}) \\ &\quad + f(z|\mathbf{\Omega}; db, db)\Pr(db, db|\mathbf{\Omega}). \end{aligned}$$

Write $\Pr(\mathbf{\Omega}|sm, db) = \Pr(\mathbf{\Omega}|sm, db, y > z)\Pr(y > z) + \Pr(\mathbf{\Omega}|sm, db, y < z)\Pr(y < z)$, where "$y > z$" means the event that the unknown value of $\tilde{Y}$ is higher than the unknown value of $\tilde{Z}$. Of course, $\Pr(y > z) = \Pr(z > y) = 1/2$. Define $LA(\bar{y}) := \{z : z \in A \wedge z < \bar{y}\}$ and $HA(\bar{y}) := \{z : z \in A \wedge z > \bar{y}\}$. Thus

$$\begin{aligned} \Pr(\mathbf{\Omega}|sm, db, y > z) &= \Pr\left(s^Y, s^Z|sm, db, y > z\right)\Pr\left(\bar{y}, z \in A|sm, db, y > z, s^Y, s^Z\right) \\ &= \frac{1}{2}\frac{\int_{LA(\bar{y})} f(\bar{y}, z)\,dz}{\Pr(y > z)} = f(\bar{y})\int_{LA(\bar{y})} f(z)\,dz, \end{aligned}$$

and analogously $\Pr(\mathbf{\Omega}|sm, db, z > y) = \Pr(\mathbf{\Omega}|db, sm, y > z) = 0$, $\Pr(\mathbf{\Omega}|db, sm, z > y) = f(\bar{y})\int_{HA(\bar{y})} f(z)\,dz$, $\Pr(\mathbf{\Omega}|db, db, y > z) = \frac{1}{2}f(\bar{y})\int_{LA(\bar{y})} f(z)\,dz$, and $\Pr(\mathbf{\Omega}|db, db, z > y) = \frac{1}{2}f(\bar{y})\int_{HA(\bar{y})} f(z)\,dz$. Thus, $\Pr(\mathbf{\Omega}|sm, db) = \frac{1}{2}f(\bar{y})\int_{LA(\bar{y})} f(z)\,dz$, $\Pr(\mathbf{\Omega}|db, sm) = \frac{1}{2}f(\bar{y})\int_{HA(\bar{y})} f(z)\,dz$ and $\Pr(\mathbf{\Omega}|db, db) = \frac{1}{4}f(\bar{y})\int_A f(z)\,dz$. Therefore, $\Pr(\mathbf{\Omega}) = \frac{1}{4}(1 - \pi^2)f(\bar{y})\int_A f(z)\,dz$. Using Bayes' rule, one finds,

$$\begin{aligned} \Pr(sm, db|\Omega) &= \frac{\frac{1}{2}f(\bar{y})\int_{LA(\bar{y})} f(z)\,dz\,(1-\pi)\pi}{\frac{1}{4}(1-\pi^2)f(\bar{y})\int_A f(z)\,dz} = \frac{2\pi}{1+\pi}\frac{\int_{LA(\bar{y})} f(z)\,dz}{\int_A f(z)\,dz}. \\ \Pr(db, sm|\Omega) &= \frac{\frac{1}{2}f(\bar{y})\int_{HA(\bar{y})} f(z)\,dz\pi(1-\pi)}{\frac{1}{4}(1-\pi^2)f(\bar{y})\int_A f(z)\,dz} = \frac{2\pi}{1+\pi}\frac{\int_{HA(\bar{y})} f(z)\,dz}{\int_A f(z)\,dz} \\ \Pr(db, db|\Omega) &= \frac{\frac{1}{4}f(\bar{y})\int_A f(z)\,dz\,(1-\pi)^2}{\frac{1}{4}(1-\pi^2)f(\bar{y})\int_A f(z)\,dz} = \frac{1-\pi}{1+\pi}. \end{aligned}$$

What remains to be determined are the densities conditional on each pair of abilities. Note that the event $\left\{s^Y, s^Z, \bar{y}, z \in A, sm, db\right\}$ implies that $z < y$ and $z \in A$. We will write this event as "$\mathbf{\Omega}; sm, db$". Thus,

$$\begin{aligned} f(z|\mathbf{\Omega}; sm, db) &= \frac{f(z|y)}{\int_{LA(y)} f(z|y)\,dz} I_{LA(y)}(z) = \frac{f(z)}{\int_{LA(y)} f(z)\,dz} I_{LA(y)}(z) \\ f(z|\mathbf{\Omega}; db, sm) &= \frac{f(z)}{\int_{HA(y)} f(z)\,dz} I_{HA(y)}(z) \\ f(z|\mathbf{\Omega}; db, db) &= \frac{f(z)}{\int_A f(z)\,dz} I_A(z). \end{aligned}$$

The result that $f(z|\boldsymbol{\Omega}) = \frac{f(z)}{\int_A f(z)dz} I_A(z)$ can now readily be obtained.■

Proof of Proposition 3: (i) We first study $X_{2,1} = Z$. Consider the decision-making stage. With messages sent and beliefs set by the market, agent 1 continues with Y if and only if

$$\begin{aligned} & y + \lambda \Pr\left(X_{2,2} = Y|s^Y, s^Z, m_1, m_2\right) \hat{\pi}_1(Y, Z, Y, Y) \\ & +\lambda \Pr\left(X_{2,2} = Z|s^Y, s^Z, m_1, m_2\right) \hat{\pi}_1(Y, Z, Y, Z) \\ \geq\ & E\left[\widetilde{Z}|s^Y, s^Z, y, m_2\right] + \lambda \Pr\left(X_{2,2} = Y|s^Y, s^Z, m_1, m_2\right) \hat{\pi}_1(Y, Z, Z, Z) \\ & +\lambda \Pr\left(X_{2,2} = Z|s^Y, s^Z, m_1, m_2\right) \hat{\pi}_1(Y, Z, Z, Y). \end{aligned} \tag{A1}$$

As $E\left[\widetilde{Z}|s^Y, s^Z, y, m_2\right] \equiv E\left[\widetilde{Z}|m_2\right]$ is independent of y, see Lemma A1, it follows that agent 1's decision strategy is a threshold strategy irrespective of the decision strategy of 2. The unique equilibrium decision strategies are therefore a pair of threshold strategies. Given the symmetry of the model, it is natural to focus on the situation in which both agents use the same threshold value. The market then draws the following inferences. If 1 keeps Y and 2 switches to Y, then $y > z$. The implication is that 1 received a correct signal while 2 received an incorrect signal, i.e., $\hat{\pi}_1(Y, Z, Y, Y) = \frac{2\pi}{1+\pi}$ and $\hat{\pi}_2(Y, Z, Y, Y) = 0$. Analogously, $\hat{\pi}_1(Y, Z, Z, Z) = 0$ and $\hat{\pi}_2(Y, Z, Z, Z) = \frac{2\pi}{1+\pi}$. In case that neither agent switches or that both agents switch, the market learns nothing about the relative value of y and z. Then, the only relevant information for the market when updating its beliefs is that in $t = 1$ the agents received different signals. From this the market infers that at least one of the agents received an uninformative signal, and thus $\hat{\pi}_1(Y, Z, Y, Z) = \hat{\pi}_1(Y, Z, Z, Y) = \frac{\pi}{1+\pi}$.

For given reputations and behaviour of 2, if 1 observes $y = \bar{y}_{\mathsf{D}}^{\mathsf{gl}}$ and continues with Y he gets $\bar{y}_{\mathsf{D}}^{\mathsf{gl}} + \lambda \Pr\left(z < \bar{y}_{\mathsf{D}}^{\mathsf{gl}}\right) 2\pi/(1+\pi) + \lambda \Pr\left(z \geq \bar{y}_{\mathsf{D}}^{\mathsf{gl}}\right) \pi/(1+\pi)$, whereas switching to Z yields $E\left[\widetilde{Z}|m_2\right] + \lambda \Pr\left(z < \bar{y}_{\mathsf{D}}^{\mathsf{gl}}\right) \pi/(1+\pi)$. 1 continues with Y if and only if

$$\bar{y}_{\mathsf{D}}^{\mathsf{gl}} \geq E\left[\widetilde{Z}|m_2\right] - \lambda \frac{\pi}{1+\pi}, \tag{A2}$$

and analogously for 2. Now consider the communication stage. As decisions are made according to (A2), the message that an agent sends does not influence his own decision. Moreover, independent of the technology an agent uses in $t = 2$, his ex post reputation is strengthened by a switch by the other agent to his initial technology, $\hat{\pi}_1(Y, Z, Y, Z) < \hat{\pi}_1(Y, Z, Y, Y)$ and $\hat{\pi}_1(Y, Z, Z, Z) < \hat{\pi}_1(Y, Z, Z, Y)$. The implication is that the interest an agent has to convince the other agent to switch technology destroys all meaningful communication. The unique equilibrium communication strategy in case $X_{2,1} = Z$ is a pooling strategy.[45] Thus, in equilibrium $\bar{y}_{\mathsf{D}}^{\mathsf{gl}}$ is determined by (8), and is thus a decreasing function of λ.

(ii) If $X_{2,1} = Y$, truthful revelation is an equilibrium strategy, but communication is also irrelevant. That the point of indifference $\bar{y}_{\mathsf{S}}^{\mathsf{gl}}$ satisfies (7) is immediate. The comparative statics result on $\bar{y}_{\mathsf{S}}^{\mathsf{gl}}$ uses Theorem 3 in Milgrom and Roberts (1994), see also proof of Proposition 1. The expressions for $\underline{\lambda}_{\mathsf{dl}}^{\mathsf{gl}}$ and $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{gl}}$ are then immediate.

We conclude by providing more details about the posteriors. $\hat{\pi}_1 (YYYY; \bar{c}) = \Pr\left(\bar{\theta}|YYYY; \bar{c}\right)$. Write $F(\bar{c}) = F$. Use $\Pr\left(YYYY|\bar{\theta}\right) = \Pr\left(YYYY|\bar{\theta}, y > z\right)/2 = (1+\pi)\Pr(y > \bar{c}|y > z)/4 = (1+\pi)(1-F^2)/4$ and $\Pr(YYYY|\underline{\theta}) = (1+\pi)(1-F^2)/8 + (1-\pi)(1-F)^2/8$, and so $\hat{\pi}(YYYY; \bar{c}) = (1+F)\frac{1+\pi}{1+\pi^2+2F\pi}\pi > \pi$. Similarly, $\hat{\pi}(YYZZ; \bar{c}) = F\frac{\pi+1}{1+\pi(2F-2+\pi)}\pi$. One can check that $\Delta\hat{\pi}(\mathsf{S}) := \hat{\pi}(YYYY; \bar{c}) - \hat{\pi}(YYZZ; \bar{c})$ is decreasing in $\bar{c}$. In particular, for $\bar{y}_{\mathsf{S}}^{\mathsf{gl}} = 0$, the gap equals $\frac{1+\pi}{1+\pi^2}\pi$. Also, $\hat{\pi}_1(YZYY) = \Pr\left(\bar{\theta}|YZYY\right)$. From $\{Y, Z, Y, Y\}$ the market deduces that $y > z$ in case of both first-best and equilibrium behaviour. Thus, $\Pr\left(YZYY|\bar{\theta}\right) = (1-\pi)/4$ (as $\theta_2 = \underline{\theta}$ for $X_{2,1} = Z$) and $\Pr(YZYY|\underline{\theta}) = (1-\pi)/8$, and apply Bayes rule. Finally, $\hat{\pi}_1(YZYZ) = \Pr\left(\bar{\theta}|YZYZ\right)$. From $\{Y, Z, Y, Z\}$ the market deduces that $(y, z) \in A := \left\{(y, z)\,|y, z < \bar{y}_{\mathsf{D}}^{\mathsf{gl}} \text{ or } y, z > \bar{y}_{\mathsf{D}}^{\mathsf{gl}}\right\}$. Use $\Pr\left(YZYZ|\bar{\theta}\right) = \frac{1-\pi}{2}\Pr(A|y > z)\frac{1}{2}$, $\Pr(YZYZ|\underline{\theta}) = \frac{1}{2}\frac{1-\pi}{2}\Pr(A|y > z)\frac{1}{2} + \frac{1}{2}\frac{1+\pi}{2}\Pr(A|z > y)\frac{1}{2}$, and $\Pr(A|z > y) = \Pr(A|y > z)$ (as Y and Z are iid), and apply Bayes rule. ■

Proof of Proposition 4: We begin by completing the statement of the proposition. Define a truncated density as follows: $Tr(x; a_r, a_{r+1}) = g(x)/(F(a_{r+1}) - F(a_r))$, where $g(x) = f(x)$ for $x \in [a_r, a_{r+1}]$ and $g(x) = 0$ everywhere else. The centre's belief function is (a) $f_C(y|\mathbf{I}_C) = Tr\left(y; a_r^*, a_{r+1}^*\right)$ for $m_1 \in \left(a_r^*, a_{r+1}^*\right)$ and $f_C(z|\mathbf{I}_C) = Tr\left(z; a_r^*, a_{r+1}^*\right)$ for $m_2 \in \left(a_r^*, a_{r+1}^*\right)$, for $r = 0, \ldots, N^* - 1$ if initial technologies differ, and (b) $f_C(y|\mathbf{I}_C) = Tr(y; 0, \bar{y}_{\mathsf{S}}^*)$ for $m_1 \in [0, \bar{y}_{\mathsf{S,cl}}]$ and $f_C(x_{1,1}|\mathbf{I}_C) = Tr(y; \bar{y}_{\mathsf{S,cl}}, 1)$ for $m_1 \in (\bar{y}_{\mathsf{S,cl}}, 1]$ if initial technologies are

[45] To avoid a discussion of out-of-equilibrium beliefs, we assume that each agent uses a probability distribution over the full support $[0, 1]$ that is independent of the value x he observed. We refer to this equilibrium communication strategy simply by "pooling strategy".

the same. In case of lo, the partition $\mathbf{a}^*$ and the collusion strategy $\bar{y}_{\mathsf{S,cl}} = \bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}}$ satisfy

$$\lambda \left[\hat{\pi}\left(Y,Y;\bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}},\mathbf{a}^*\right) - \hat{\pi}\left(Y,Z;\bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}},\mathbf{a}^*\right)\right] = E\left[\widetilde{Z}|a^*_{r-1} \leq z \leq a^*_{r+1}\right] - a^*_r \quad \text{(A3)}$$

$$\lambda \left[\hat{\pi}\left(Y,Y;\bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}},\mathbf{a}^*\right) - \hat{\pi}\left(Y,Z;\bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}},\mathbf{a}^*\right)\right] = E\left[\widetilde{Z}|s^Y,s^Y,\bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}}\right] - \bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}} \quad \text{(A4)}$$

for $r = 1,\ldots,N^* - 1$. In case of gl, the partition $\mathbf{a}^*$ satisfies

$$\lambda\left[\hat{\pi}_1\left(Y,Z,Y,Y;\mathbf{a}^*\right) - \hat{\pi}_1\left(Y,Z,Z,Z;\mathbf{a}^*\right)\right] = E\left[\widetilde{Z}|a^*_{r-1} \leq z \leq a^*_{r+1}\right] - a^*_r \quad \text{(A5)}$$

for $r = 1,\ldots,N^* - 1$. The collusion strategy $\bar{y}_{\mathsf{S,cl}} = \bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}}$ satisfies

$$\lambda\left[\hat{\pi}\left(Y,Y,Y,Y;\bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}}\right) - \hat{\pi}\left(Y,Y,Z,Z;\bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}}\right)\right] = E\left[\widetilde{Z}|s^Y,s^Y,\bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}}\right] - \bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}}. \quad \text{(A6)}$$

(A) Consider $X_{2,1} = Z$. Assume that the centre uses (9). As $\mathbf{I}^{\mathsf{cl}}_C = \left\{s^Y,s^Z,m^Y,m^Z\right\}$, $E\left[\tilde{Y}|\mathbf{I}^{\mathsf{cl}}_C\right] = E\left[\tilde{Y}|m^Y\right]$ and $E\left[\tilde{Z}|\mathbf{I}^{\mathsf{cl}}_C\right] = E\left[\tilde{Z}|m^Z\right]$. We start by showing that irrespective of the communication strategy of agent 2, agent 1 uses a partition strategy. We are done if we have shown that (i) separation in an interval (y',y'') cannot happen and (ii) that the set of values of y for which 1 sends a given message is convex. Note that as a result of the centre's objective function her decision rule is such that whenever $E\left[\tilde{Y}|m^Y\right] \neq E\left[\tilde{Z}|m^Z\right]$ she imposes the better technology at both sites. Hence, if an agent is allowed to continue with his technology, this can not hurt his reputation: neither $\hat{\pi}_1(YY) < \hat{\pi}_1(YZ)$ nor $\hat{\pi}_1(YZYY) < \hat{\pi}_1(YZZZ)$ can hold in equilibrium. In what follows we limit attention to reputations that are locally determined as the proof for globally determined reputations proceeds analogously.

Ad (i). We start by assuming that agent 1 separates for all $y \in (y',y'')$ and derive a contradiction. If 1 separates in $y \in (y',y'')$, then $\hat{\pi}_1(YY) > \hat{\pi}_1(YZ)$. To see that $\hat{\pi}_1(YY) = \hat{\pi}_1(YZ)$ cannot hold, note that this equality would require that for *all* pairs of messages that the centre receives, $E\left[\tilde{Y}|m^Y\right] = E\left[\tilde{Z}|m^Z\right]$. And this, in turn, would require in particular that agent 1 does not separate for any value of y, a contradiction. Thus, if agent 1 separates for all $y \in (y',y'')$, then $\hat{\pi}_1(YY) > \hat{\pi}_1(YZ)$. Now suppose 1 observes some $y \in (y',y'')$. Rather than telling the truth, he prefers to exaggerate and send a message $m = y + \varepsilon$, with $0 < \varepsilon < \lambda(\hat{\pi}_1(YY) - \hat{\pi}_1(YZ))$. Then, conditional on the exaggeration inducing the centre to impose Y rather than Z at both sites, his net change in utility is positive. That is, truthtelling cannot be part of an equilibrium.

Ad (ii). To show that the set of values of y for which a given message is sent is convex, it suffices to show that if 1 sends m_1^Y both for y_1 and $y_2 > y_1$, then he also sends it for any $y' \in (y_1, y_2)$. Define the sets $H(m^Y) = \left\{z : E\left[\tilde{Y}|m^Y\right] > E\left[\tilde{Z}|m^Z(z)\right]\right\}$, $G(m^Y) = \left\{z : E\left[\tilde{Y}|m^Y\right] = E\left[\tilde{Z}|m^Z(z)\right]\right\}$, and $L(m^Y) = \left\{z : E\left[\tilde{Y}|m^Y\right] < E\left[\tilde{Z}|m^Z(z)\right]\right\}$, and the probabilities $p(m^Y) := \int_{H(m^Y)} f(z)\,dz$ and $q(m^Y) := \int_{G(m^Y)} f(z)\,dz$. Hence, if 1 sends message m_1^Y if he has observed y, then in equilibrium his expected payoffs equal

$$[y + \lambda\hat{\pi}_1(Y)]\left(p(m_1^Y) + \frac{q(m^Y)}{2}\right) + \frac{1}{2}\int_{G(m^Y)} [z + \lambda\hat{\pi}_1(Z)] f(z)\,dz + \int_{L(m^Y)} [z + \lambda\hat{\pi}_1(Z)] f(z)\,dz,$$

and must be at least as large as the expected payoff from sending $m_2^Y \neq m_1^Y$. Note that the expression is linear in y. It then follows that if agent 1 sends m_1^Y both for y_1 and $y_2 > y_1$, then he also sends it for any $y' \in (y_1, y_2)$.

So far, we have established that each agent uses a partition strategy. Assume now that reputations are given, and that agent 2 uses the partition strategy $(N^*, \mathbf{a}^*)$ to communicate about Z. We show that it is then a best-reply for agent 1 to use a partition strategy with the same partitions to communicate about Y. Write $\hat{\pi}(Y, X)$ instead of $\hat{\pi}\left(Y, X; \bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}}, \mathbf{a}^*\right)$. Let $y = a_r$, where we have suppressed reference to the number of partitions N. At this value of y, 1 should be indifferent between sending some $m_{r+1} \in [a_r, a_{r+1})$ or some $m_r \in [a_{r-1}, a_r)$. If $z < a_{r-1}$ or $z \geq a_{r+1}$, whether 1 sends m_r or m_{r+1} does not affect the decision of the centre. Hence, one can limit attention to $z \in [a_{r-1}, a_{r+1})$. As $E\left[\widetilde{Z}|s^Y, s^Z, y = a_r\right] = E\left[\widetilde{Z}\right]$, $E\left[\widetilde{Z}|s^Y, s^Z, y = a_r, \alpha \leq z \leq \beta\right] = E\left[\widetilde{Z}|\alpha \leq z \leq \beta\right]$ for any pair (α, β) such that $0 \leq \alpha < \beta \leq 1$. Let $p(\alpha, \beta) := F(\beta) - F(\alpha)$. Sending m_{r+1} yields agent 1

$$\begin{aligned} &p(a_{r-1}, a_r)[a_r + \lambda\hat{\pi}_1(YY)] + \frac{1}{2}p(a_r, a_{r+1})[a_r + \lambda\hat{\pi}_1(YY)] + \\ &\frac{1}{2}p(a_r, a_{r+1})\left[E\left[\widetilde{Z}|a_r \leq z < a_{r+1}\right] + \lambda\hat{\pi}_1(YZ)\right], \end{aligned} \tag{A7}$$

whereas m_r yields

$$\begin{aligned} &\frac{1}{2}p(a_{r-1}, a_r)[a_r + \lambda\hat{\pi}_1(YY)] + \frac{1}{2}p(a_{r-1}, a_r)\left[E\left[\widetilde{Z}|a_{r-1} \leq z < a_r\right] + \lambda\hat{\pi}_1(YZ)\right] \\ &+p(a_r, a_{r+1})\left[E\left[\widetilde{Z}|a_r \leq z < a_{r+1}\right] + \lambda\hat{\pi}_1(YZ)\right]. \end{aligned} \tag{A8}$$

Equating (A7) and (A8) shows that agent 1 is indifferent between sending m_{r+1} and m_r for $y = a_r$ if (A3) holds.

(B) If $X_{2,1}=Y$, it is straightforward to check that, if agent 2 uses the collusion strategy, if the centre's decision strategy is as stated, and for given beliefs $\hat{\pi}$, then for agent 1 a collusion strategy with $\bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}}$ satisfying (A4) is a best-reply. It is straightforward to establish that the belief function follows from applying Bayes' rule to the communication strategies of the agents, and that the centre's decision strategy is a best reply given the belief function.

Consider $\bar{y}^{\mathsf{lo}}_{\mathsf{S,cl}}=0, \mathbf{a}^*_1(1)=0$ (a pooling communication strategy). To determine $\hat{\pi}(YY)-\hat{\pi}(YZ)$, use $\Pr\left(YY|\bar{\theta}\right)=\Pr\left(YY|\bar{\theta},y>z\right)\frac{1}{2}+\Pr\left(YY|\bar{\theta},z>y\right)\frac{1}{2}=\frac{1}{8}\pi+\frac{3}{8}$, and $\Pr\left(YY|\underline{\theta}\right)=\frac{3}{8}$. Hence, $\hat{\pi}(YY)=\frac{3+\pi}{3+\pi^2}\pi$. Similarly, $\hat{\pi}(YZ)=\frac{\pi}{\pi+1}$, such that $\hat{\pi}(YY)-\hat{\pi}(YZ)=\frac{4\pi}{(3+\pi^2)(1+\pi)}\pi$. The RHS of both (A3) and (A4) become $E\left[\widetilde{Z}\right]$. Hence, this communication strategy is indeed the equilibrium for $\lambda\geq\bar{\lambda}^{\mathsf{lo}}_{\mathsf{cl}}$.

Now turn to gl. Assume $X_{1,1}=X_{2,1}=Y$. For given parameter values the collusion strategy is the same as the cut-off strategy in case of dl cum gl. Thus, $\hat{\pi}\left(YYYY;\bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}}=0\right)=(1+F(0))\frac{1+\pi}{1+\pi^2+2F(0)\pi}\pi=\pi(1+\pi)/(1+\pi^2)$, and $\hat{\pi}\left(YYZZ;\bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}}=0\right)=0$, and the RHS of (A6) becomes $E\left[\widetilde{Z}\right]$ for $\bar{y}^{\mathsf{gl}}_{\mathsf{S,cl}}=0$. Hence, this collusion strategy is indeed the equilibrium strategy for $\lambda\geq\bar{\lambda}^{\mathsf{gl}}_{\mathsf{cl}}$. Now assume $X_{1,1}\neq X_{2,1}$. Assume $N=2$, and define $a:=a_1$. $\hat{\pi}_1(YZYY;a)=\Pr\left(\bar{\theta}|YZYY;a\right)$. Use

$$\begin{aligned}\Pr\left(YZYY|\bar{\theta};a\right) &= \frac{1}{2}\frac{1-\pi}{2}\left(\Pr(y>a>z|y>z)+\Pr(y>z>a|y>z)\frac{1}{2}+\Pr(a>y>z|y>z)\frac{1}{2}\right)\\ &= \frac{1-\pi}{4}\left(\frac{1}{2}+F(a)(1-F(a))\right).\end{aligned}$$

Similarly, $\Pr\left(YZYY|\bar{\theta};a\right)=\frac{1}{8}-\frac{\pi}{4}F(a)(1-F(a))$. Hence, $\hat{\pi}_1(YZYY)=\frac{\pi}{1+\pi}\left(1+2F(a)-2F(a)^2\right)$. Analogously, $\hat{\pi}_1(YZZZ)=\frac{\pi}{1+\pi}\left(1-2F(a)+2F(a)^2\right)$, and the reputational gap becomes $4\frac{\pi}{1+\pi}F(a)(1-F(a))$. As $\lambda\frac{4\pi}{1+\pi}F(0)(1-F(0))=0<E\left[\widetilde{Z}\right]$ and $\lambda\frac{4\pi}{1+\pi}F\left(E\left(\widetilde{Z}\right)\right)\left(1-F\left(E\left[\widetilde{Z}\right]\right)\right)>0$, for all continuous F and any finite λ there is a unique $a^*_1>0$ that satisfies (A5). That is, for any finite λ, $N^*\geq 2$.■

Proof of Proposition 5: The existence of λ_1 and the role played by the inequality $E\left[\widetilde{Z}\right](1+\pi)<1$ in the existence of λ_2 and λ_3 was shown in the text preceding the proposition. Numerical calculations show that for $f_X=1$, $W_{\mathsf{ia}}(\lambda,\pi)<W^{\mathsf{gl}}_{\mathsf{dl}}(\lambda,\pi)<W^{\mathsf{lo}}_{\mathsf{dl}}(\lambda,\pi)$ holds for all λ and π.■

Proof of Proposition 6: Partly shown in the text preceding the proposition. Concerning the relative performance of ia, suppose $\lambda=0$. Then, $W_{\mathsf{ia}}(0,\pi)$ is equal to $W(0,\pi)$ in case *one* agent reports to the centre that $y \geq \bar{y}_{\mathsf{ia}}^{FB}$ or $y < \bar{y}_{\mathsf{ia}}^{FB}$. In case of cl, *two* agents reveal information truthfully to the centre. By continuity of $W_{\mathsf{ia}}(\lambda, \pi)$ and $W_{\mathsf{cl}}(\lambda, \pi)$ in λ, $W_{\mathsf{ia}}(\lambda, \pi) < W_{\mathsf{cl}}(\lambda, \pi)$ for all $\lambda < \lambda_4$, for some $\lambda_4 > 0$. The second part of the proposition follows from the facts that (i) $\bar{\lambda}_{\mathsf{cl}}^{\mathsf{lo}} > \bar{\lambda}_{\mathsf{ia}}$ (see Propositions 1 and 4), and (ii) for all λ agents send influential information under cl *cum* gl.■

Proof of Proposition 7: Fix λ, π, and f_X. Suppose $X_{1,1} = X_{2,1}$. A straightforward comparison of (7) and (A6) shows that welfare is the same under dl and cl for all f_X, π, and λ. Now suppose $X_{1,1} \neq X_{2,1}$. In case of cl and in equilibrium, the more ranks the agents use, the higher is W. Hence, it suffices to show that the proposition is true if communication under cl is limited to two ranks. Proposition 4 *(iv)* shows that an equilibrium with two ranks exists for all parameter values. This partition is characterized by $a_1^* \in \left(0, E\left[\widetilde{Z}\right]\right)$. Thus, if agents rank their technologies differently, the centre picks the higher ranked technology. Given the communication strategies of the agents this technology is indeed the better one. However, for $(y, z) \in [0, a_1^*]^2$ and $(y, z) \in [a_1^*, 1]^2$, both technologies are ranked in the same way. Hence, the centre tosses a fair coin. The inferior technology is chosen half of the time at both sites. In case of dl, for $y < \bar{y}_{\mathsf{D}}^{\mathsf{gl}} \leq z$, the Y-user switches to Z, and the Z-user continues his technology. Both agents use the superior technology in $t = 2$. The same holds, mutatis mutandis, for $z < \bar{y}_{\mathsf{D}}^{\mathsf{gl}} \leq y$. However, for $(y, z) \in \left[0, \bar{y}_{\mathsf{D}}^{\mathsf{gl}}\right]^2$, both agents switch, while if $(y, z) \in \left[\bar{y}_{\mathsf{D}}^{\mathsf{gl}}, 1\right]^2$, both agents continue. In either case, the inferior technology is used at one site with probability one. Clearly, if $a_1^* = \bar{y}_{\mathsf{D}}^{\mathsf{gl}}$, then cl and dl would yield the same expected welfare. For given parameter values, they are, however, not the same. $\bar{y}_{\mathsf{D}}^{\mathsf{gl}}$ satisfies $\lambda \frac{\pi}{1+\pi} = E\left[\widetilde{Z}\right] - \bar{y}_{\mathsf{D}}^{\mathsf{gl}}$ (see (8)), whereas a_1^* satisfies $\lambda \frac{\pi}{1+\pi} 4F(a_1^*)(1 - F(a_1^*)) = E\left[\widetilde{Z}\right] - a_1^*$ (see (10)). As $4F(a_1)(1 - F(a_1)) < 1$ for all a_1, for given parameter values, the reputational gap in case of cl is smaller than in case of dl. As this gap equals the size of the distortion, $E\left[\widetilde{Z}\right] - a_1^*$ or $E\left[\widetilde{Z}\right] - \bar{y}_{\mathsf{D}}^{\mathsf{gl}}$, cl yields a higher expected welfare than dl.■

References

Austen-Smith, David. 1990. "Information Transmission in Debate." *American Journal of Political Science* 34 (1): 124-152.

Alonso, Ricardo, Wouter Dessein and Niko Matouschek. 2008. "When Does Coordination Require Centralization?" *American Economic Review* 98 (1): 145-179.

Austen-Smith, David, and Timothy Feddersen. 2005. "Deliberation and Voting Rules." In *Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks*, edited by David Austen-Smith and John Duggan, 269–316. Berlin: Springer-Verlag.

Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning from Neighbours." *Review of Economic Studies* 65 (3): 595–621.

Bewley, Susan. 2008. "Getting to the bottom of evidence based medicine." *British Medical Journal* 336 (7647): 764.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1998. "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades." *Journal of Economic Perspectives* 12 (3): 151-170.

Bergemann, Dirk and Juuso Välimäki. 2006. "Bandit Problems." Cowles Foundation Discussion Paper 1551.

Çelen, Boğaçhan, Shachar Kariv, and Andrew Schotter. 2010. "An Experimental Test of Advice and Social Learning." *Management Science* 56 (10): 1678-1701.

Chen, Ying, Navin Kartik, and Joel Sobel. 2008. "Selecting Cheap-Talk Equilibria." *Econometrica* 76 (1): 117-136.

Coughlan, Peter J. 2000. "In Defense of Unanimous Jury Verdicts: Mistrials, Communication, and Strategic Voting." *American Political Science Review* 94 (2): 375–393.

Dessein, Wouter. 2007. "Why a Group Needs a Leader: Decision-making and Debate in Committees." Unpublished manuscript.

Eddy, David M. 1990. "Clinical decision-making: From Theory to Practice. The Challenge." *Journal of the American Medical Association* 263 (2): 287-290.

Farrell, Joe. 1993. "Meaning and Credibility in Cheap-Talk Games." *Games and Economic Behavior* 5 (4): 514-531.

Fink, Arlene, Jacqueline Kosecoff, Mark Chassin, and Robert H. Brook 1984. "Consensus Methods: Characteristics and Guidelines for Use." *American Journal of Public Health* 74 (9): 979-983.

Fisher, Paul E., and Philip C. Stocken. 2001. "Imperfect Information and Credible Communication." *Journal of Accounting Research* 39 (1): 119-134.

Friebel, Guido and Michael Raith. 2010. "Resource Allocation and Firm Scope." *American Economic Journal: Microeconomics* 2 (2): 1-33.

Gilligan, Thomas W. and Keith Krehbiel. 1987. "Collective Decisionmaking and Standing Committees: an Informational Rationale for Restrictive Amendment Procedures." *Journal of Law, Economics, and Organization* 3 (2): 287-335.

Gibbons, Robert, and Rebecca Henderson. 2013. "What Do Managers Do? Exploring Persistent Performance Differences among Seemingly Similar Enterprises." In *The Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, 680-731. Princeton University Press, Princeton.

Ivanov, Maxim. 2010. "Informational Control and Organizational Design." *Journal of Economic Theory* 145 (2): 721-751.

Kanodia, Chandra, Robert Bushman, and John Dickhaut. "Escalation Errors and the Sunk Cost Effect: An Explanation Based on Reputation and Information Asymmetries." *Journal of Accounting Research* 27 (1): 59-77.

Kartik, Navin. 2009. "Strategic Communication with Lying Costs." Review of Economic Studies 76 (4): 1359-1395.

Li, Zhouzheng, and Huanxing Yang. 2013. "Competitive Cheap Talk." Unpublished manuscript.

Milbourn, Todd T., Shockley, Richard L., Thakor, Anjan V. 2001. "Managerial career concerns and investments in information." *RAND Journal of Economics* 32 (2): 334– 351.

Milgrom, Paul, and John Roberts. 1986. "Relying on the Information of Interested Parties." *Rand Journal of Economics* 17 (1): 18-32.

Milgrom, Paul, and John Roberts. 1994. "Comparing Equilibria." *American Economic Review* 84 (3): 441-459.

Mookherjee, Dilip. 2006. "Decentralization, Hierarchies and Incentives: A Mechanism Design Perspective." *Journal of Economic Literature* 44 (2): 367-390.

Oates, Wallace E. 1999. "An Essay on Fiscal Federalism." *Journal of Economic Literature* 37 (3): 1120-1149.

Ottaviani, Marco and Peter Norman Sørensen. 2001. "Information aggregation in debate:

who should speak first?" *Journal of Public Economics* 81 (3): 393-421.

Ottaviani, Marco and Peter Norman Sørensen. 2006. "Reputational Cheap Talk." *Rand Journal of Economics* 37 (1): 155-175.

Phelps, Charles E. 2000. "Information Diffusion and Best Practice Adoption." In *Handbook of Health Economics, Volume 1A*, edited by Anthony J. Culyer and Joseph P. Newhouse, 223-264. Amsterdam: Elsevier, Amsterdam.

Pochet, Phillipe. 2005. "The Open Method of Co-ordination and the Construction of Social Europe. A Historical Perspective." In *The Open Method of Co-ordination in Action. The European Employment and Social Inclusion Strategies*, edited by Jonathan Zeitlin and Phillipe Pochet, 37-82. Brussels: PIE-Peter Lang.

Poitevin, Michel. 2000. "Can the Theory of Incentives Explain Decentralization?" *Canadian Journal of Economics* 33 (4): 878-906

Prat, Andrea. 2005. "The Wrong Kind of Transparency." *American Economic Review* 95 (3): 862-877.

Prendergast, Canice and Stole. 1996. "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning." *Journal of Political Economy* 104 (6): 1105-1134.

Radaelli, Claudio M. 2003. "The Open Method of Coordination: a New Governance Architecture for the European Union?" SIEPS report 2003-1.

Rantakari, Heikki. 2008. "Governing Adaptation." *Review of Economic Studies* 75 (4): 1257-1285.

Rantakari, Heikki. 2014. "A Simple Model of Project Selection With Strategic Communication and Uncertain Motives." *Journal of Economic Behavior and Organization* 102 (6): 14-42.

Rowe, Gene, George Wright, and Fergus Bolger. 1991. "Delphi: A Reevaluation of Research and Theory." *Technological Forecasting and Social Change* 39 (3): 235-251.

Smith, Lones and Peter Sørensen. 2000. "Pathological Outcomes of Observational Learning." *Econometrica* 68 (2): 371-398.

Suurmond, Guido, Otto H. Swank and Bauke Visser. 2004. "On the Bad Reputation of Reputational Concerns." *Journal of Public Economics* 88 (12): 2817-2838.

Swank, Otto H. and Bauke Visser. 2010. "Decision Making and Learning in a Globalizing

World." Tinbergen Institute Discussion Paper 10-034/1.

Swank, Otto H. and Bauke Visser. 2013. "Is Transparency to no Avail? Committee Decision-Making, Pre-meetings, and Credible Deals." *Scandinavian Journal of Economics* 115 (4): 967-994.

Thaler, Richard H. 1980. "Towards a Positive Theory of Consumer Choice." *Journal of Economic Behavior and Organization* 1 (1): 39-60.

Visser, Bauke. 1999. "Endogenous Local Interaction and Multi-Product Firms." Journal of Evolutionary Economics 9 (2): 243-263.

Visser, Bauke and Otto H. Swank. 2007. "On Committees of Experts." *Quarterly Journal of Economics* 122 (1): 337-372.

Volden, Craig, Michael M. Ting and Daniel P. Carpenter. 2008. "A Formal Model of Learning and Policy Diffusion." *American Political Science Review* 102 (3): 319-332.