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

Otto H. Swank and Bauke Visser

Erasmus University Rotterdam and Tinbergen Institute

April 2013

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. We derive a simple condition when it does. Centralizing decision rights in the presence of local markets may hurt communication so much that welfare goes down. 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.

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

*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

"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]

1 Introduction

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 et al. (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] Susan Bewley, consultant obstetrician, in Getting to the bottom of evidence based medicine, the British Medical Journal, April 5, 2008.

[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 benificial to make use of their experience. This requires communication about locally gained experience. Inspired by real world examples that we discuss below, 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 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 (*locally determined reputations*), or on the actions taken at all sites (*globally determined reputations*). In the latter case, comparisons across sites become 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 are the decision rights in period two. As a result, communication about the privately gained experience amounts to cheap talk.

We obtain the following results. In case of decentralized decision making, the quality of information exchange is high if reputations are locally determined, but low if comparisons across agents are possible. Indeed, if reputations are determined globally, an agent's reputation is particularly strong if others start to adopt "his" initial course of action. As a result, communication becomes strategic: it becomes important for an agent to convince others that "his" action is best, and communication breaks down completely. What consequence does this difference in the willingness to share information have on the quality of the decision taken? Surprisingy, less information does not necessarily mean worse decisions. This is also determined by the reputational gain of biasing the decision in favour of the action chosen in the first period. Essentially, if the agent has less information on which to base his decision, there is also less information about his ability that can be gleaned from one decision or the other, and thus less reason to distort the decision. We derive an intuitive condition that specifies when the additional information shared among agents leads to an increase or decrease in welfare.

Second period decision-making in case of a *centralized* process does not suffer from conservatism as, by assumption, the centre only cares about welfare. But the centre depends on the agents to provide him with information. An agent now faces a trade-off. On the one hand, as the agent has no 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. In equilibrium, each agent sends coarse information about his own practice. This has a number of consequences. First, replacing a decentralized process by a centralized one reduces the quality of information exchange if reputations are locally determined. We derive the conditions under which the quality of information exchange in case of a centralized process becomes so poor that it offsets the improved decision-making *conditional* on information, and lowers welfare. Second, if reputations are globally determined, replacing a decentralized process by a centralized one improves communication. As a result, a centralized process unambiguously creates more welfare than a decentralized one in case reputations are based on comparisons across sites. Furthermore, we show that vertical communication remains informative for any finite degree of conflict among agents and with the centre. Finally, we derive conditions under which, in case of centralized process, welfare goes up if locally based reputations give way to globally based reputations.

The paper is organized as follows. The next section discusses the related literature. In Section 3, we present the model. Section 4 analyses isolated agents, a benchmark situation in which agents can learn from their own past experience only. In section 5 we analyse decentralized learning, with local and global markets. In section 6 we perform the same analysis for centralized learning. Section 7 contains the comparisons. Section 8 concludes.

2 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 market circumstances and from coordinating its decision with those of the other divisions. Divisions are privately informed about their market circumstances. They can either exchange information and next decide independently of each other what decisions to take or they can report information to headquarters which then decides for both 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.

[3] Friebel and Raith (2010) study how the scope of the firm affects the quality of strategic information

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).

The importance of dispersed information has already been highlighted in debates on the relative merits of a planning economy and a market economy in the 1930s, see e.g. Hayek (1945). Team theory, as developed by Marshak and Radner (1972), is one of the first formal attempts to address the question how an organization should be structured to deal optimally with dispersed information. In this theory, interests of organizational members are perfectly aligned, and so incentive problems do not arise. Instead, 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.[4]

There are other papers in economics and political science that explore how characteristics of decision-making processes influence the quality of cheap talk communication.[5]. 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.

---

transmission between a division and head quarters.

[4]See Mookherjee (2006) and Poitevin (2000) for excellent surveys of the assumptions underlying the Revelation Principle. They also discuss various modelling strategies that can be used to explain why decentralization and delegation outperform centralization.

[5]In economics, see e.g. Dessein (2002, 2007), Visser and Swank (2007), Alonso et al. (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).

(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.[6] Part of this expert literature looks at the effects of information disclosure ('transparency') about an expert's actions and about the outcomes of decisions.[7] The present paper is related to that literature, as the information on which an agent's reputation is based can change, either by 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**.[8] In an interesting recent paper, 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.

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

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

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

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 some existing literature on learning from others. This literature is, however, methodologically quite different from ours. In the existing literature, it is *assumed* that either an agent observes the true value of the actions taken by others, whether the environment is strategic[9] or not[10], or that no such information is observed at all[11]. Furthermore, inertia or conservatism is an *exogenous* factor. For example, in the literature on word-of-mouth communication, it is assumed that only a given fraction of agents updates its decisions once new information becomes available. In our paper both 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 are facing in our model, that of choosing one 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.[12] 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. This stems from the fact that in our model the initial signal an agent receives provides information about the better technology. The higher is the observed value of a technology Y, the higher is the probability that the agent identified the better technology. And this means that it becomes more likely that the value of the other technology is lower than the

[9] 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.

[10] 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.

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

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

actual value of Y.

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.

(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.

3 A model of learning from own experience and learning from others with reputational concerns.

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 technology adopted at site i in period t is denoted by $X_{i,t}$. 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 $X_{i,t} = X$. As strategies will be defined in terms of X (or x), it will be useful to let X^C (or x^C) refer to "the other technology". That is, if $X = Y$, then $X^C = Z$, etc.

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. The informativeness of the signal depends on the agent's ability: $\Pr(s^X | x > x^C, \bar{\theta}) = 1$, $\Pr(s^X | x^C > x, \bar{\theta}) = 0$, $\Pr(s^X | x > x^C, \underline{\theta}) = \Pr(s^X | x^C > x, \underline{\theta}) = 1/2$, for $X \in \{Y, Z\}$. That is, if i is highly able, $\theta_i = \bar{\theta}$, the signal reveals with probability one the better technology: $\Pr(x > x^C | s^X, \bar{\theta}) = 1$ 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, $F_{\widetilde{X}}(x | s^X, \bar{\theta}) = F(x)^2$. On the other hand, if i is less able, $\theta_i = \underline{\theta}$, the signal is uninformative about the relative quality of the technology: $F_{\widetilde{X}}(x | s^X, \underline{\theta}) = F(x)$.[13] 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 $I_i^{\mathsf{p}} \in \mathcal{I}_i^{\mathsf{p}}$ be the information this person has at the beginning of the decision-making

[13] 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 $\theta_i = \underline{\theta}$.

stage. It depends on the process p. The decision strategy d_i^{p} determines the relationship between I_i^{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, $\mathcal{I}_i^{\mathsf{ia}} = \left\{s^Y, s^Z\right\} \times [0,1]$: 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 $d_i^{\mathsf{ia}}\left(s_i, x_{i,1}\right) \in \{Y, Z\}$ denote the technology that i uses in $t=2$ as a function of 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, $\mathcal{I}_i^{\mathsf{dl}} = \left\{s^Y, s^Z\right\} \times [0,1] \times M \times \{Y, Z\} \times M$. 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 $d_i^{\mathsf{dl}}\left(s_i, x_{i,1}, m_i, X_{j,1}, m_j\right) \in \{Y, Z\}$ denote the technology that i adopts in $t=2$ given I_i^{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, $\mathcal{I}_C^{\mathsf{cl}} = \{Y, Z\}^2 \times M^2$ 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 $d_C^{\mathsf{cl}}\left(X_{1,1}, X_{2,1}, m_1, m_2\right) \in \{Y, Z\} \times \{Y, Z\}$ 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 $\mathcal{I}_C$ instead of $\mathcal{I}_C^{\mathsf{cl}}$, 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 $\Omega_{i,t}$. We will say that "the market" infers an agent's reputation from $\Omega_{i,t}$. 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 reputations are *locally determined* if the reputation of agent i is based on the technologies used at site i only, $\Omega_{i,1} = \{X_{i,1}\}$ and

$\Omega_{i,2} = \{X_{i,1}, X_{i,2}\}$ for $i \in \{1, 2\}$. Instead, say that reputations are *globally determined* if the reputation of agent i is based on the technologies used at both sites i and j, $\Omega_{i,1} = \{X_{i,1}, X_{j,1}\}$ and $\Omega_{i,2} = \{X_{i,1}, X_{j,1}, X_{i,2}, X_{j,2}\}$ 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 $x + \lambda\hat{\pi}_{i,t}(\Omega_{i,t})$, where $\hat{\pi}_{i,t}(\Omega_{i,t}) = \Pr\left(\theta_i = \bar{\theta}|\Omega_{i,t}\right)$ equals the belief that i is highly able conditional on $\Omega_{i,t}$, 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 8.2, 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 $f_i(\cdot|I)$ 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. Because of the inherent symmetry, we write the analysis from the point of view of agent $i = 1$ and assume that $s_1 = s^Y$. Of course, $s_2 \in \left\{s^Y, s^Z\right\}$. We ignore babbling equilibria if an equilibrium in which information is transmitted exists.

4 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 $I_1^{\mathsf{ia}} = \left\{s^Y, y\right\}$, 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; panel (b) the equilibrium threshold value $\bar{y}^*_{\mathsf{ia}}$ for $\lambda < \bar{\lambda}_{\mathsf{ia}}$; panel (c) 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.

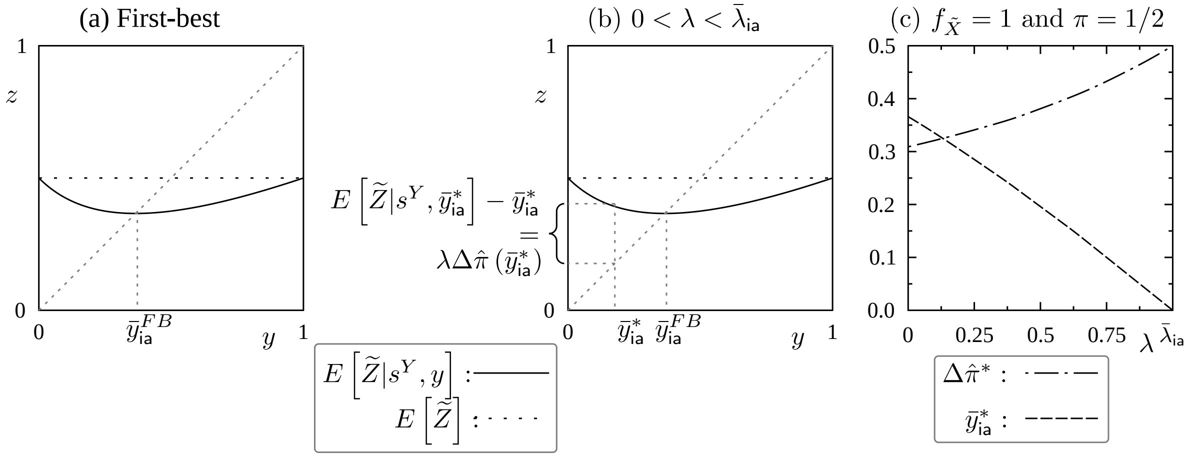

It follows from lemma 1, part (b), and it is clear from Figure 1, panel (a), that the first-best decision strategy is a *single-threshold strategy*,

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

with $\bar{t} = \bar{y}^{FB}_{\mathsf{ia}}$ 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 $\hat{\pi}\left(Y, X_{1,2}; \bar{t}\right)$ denote the reputation, obtained using Bayes' rule, for $X_{1,2} \in \{Y, Z\}$, and if the agent uses the threshold $\bar{t}$. Then,[14]

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

Irrespective of $\bar{t}$, 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 a less able agent. Hence, as an agent cares about his reputation, he wants to deviate from the

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

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 $\hat{\pi}_1(Y, Y; \bar{t}) - \hat{\pi}_1(Y, Z; \bar{t})$ 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*

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

with $\bar{y}^*_{\mathsf{ia}} \in \left(0, \bar{y}^{FB}_{\mathsf{ia}}\right)$*.* $\bar{y}^*_{\mathsf{ia}}$ *is a decreasing function of* λ*.*[15] *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. This can also be seen by rewriting (3) as $\bar{y}^*_{\mathsf{ia}} + \lambda\hat{\pi}_1(Y, Y; \bar{y}^*_{\mathsf{ia}}) = E\left[\widetilde{Z}|s^Y, \bar{y}^*_{\mathsf{ia}}\right] + \lambda\hat{\pi}_1(Y, Z; \bar{y}^*_{\mathsf{ia}})$. The left-hand side equals the value of continuing with Y if its observed value equals $\bar{y}^*_{\mathsf{ia}}$, whereas the left-hand side equals the value of switching technology for the same observed value of Y. It follows from (2) that the lower $\bar{y}^*_{\mathsf{ia}}$ is, the lower is the reputation the agent commands in case of sticking to the original technology *and* in case of switching technologies. 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. It can be checked that the reputational gap increases the lower is $\bar{y}^*_{\mathsf{ia}}$. As the reputational gap is still strictly positive for a threshold value equal to zero, it follows from (3) that for $\lambda \geq \bar{\lambda}_{\mathsf{ia}}$ $\bar{y}^*_{\mathsf{ia}} = 0$: the agent will continue with his initial choice of technology irrespective of its observed value. 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 $\hat{\pi}_1(Y, Y; \bar{y}^*_{\mathsf{ia}}) - \hat{\pi}_1(Y, Z; \bar{y}^*_{\mathsf{ia}})$, denoted by $\Delta\hat{\pi}^*$ to save space.

[15] 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, in this and all other sections.

5 Decentralized decision making

In this section, we assume that the right to decide about the technology to be adopted in period two remains with the agents. We begin by describing first-best behaviour in a decentralized process. In the communication stage each agent truthfully reveals his private information. 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. Next, the first-best decision strategy equals

$$d_1^{\mathsf{dl}}\left(I_1^{\mathsf{dl}}; \bar{y}_{\mathsf{S}}^{FB}\right) = \begin{cases} Y & \text{if } X_{2,1} = Y \text{ and } y \geq \bar{y}_{\mathsf{S}}^{FB} \\ Y & \text{if } X_{2,1} = Z \text{ and } y \geq z \\ Z & \text{otherwise,} \end{cases}$$

where $\bar{y}_{\mathsf{S}}^{FB}$ satisfies $\bar{y}_{\mathsf{S}}^{FB} = E\left[\widetilde{Z} | s^Y, s^Y, \bar{y}_{\mathsf{S}}^{FB}\right]$. That is, if both agents adopted the *same* technology, each agent should continue this technology if its value is larger than $\bar{y}_{\mathsf{S}}^{FB}$.[16] If instead agents adopted different technologies, they should next choose the one with superior performance. In 5.1 we study equilibrium behaviour in case reputations are locally determined, and in 5.2 we turn to reputations that are globally determined. In 7.1, we compare the performance of decentralized learning under both types of reputation formation.

5.1 Locally determined reputations

With agent 1's reputation independent of what agent 2 decides, his communication strategy is payoff irrelevant. 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. Let a *double-threshold strategy* $d_1^{\mathsf{dl}}\left(I_1^{\mathsf{dl}}; \bar{t}_{\mathsf{S}}, \bar{t}_{\mathsf{D}}\right)$ with thresholds $\left(\bar{t}_{\mathsf{S}}, \bar{t}_{\mathsf{D}}\right) \geqq 0$

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

be defined as

$$d_1^{\mathsf{dl}}\left(I_1^{\mathsf{dl}};\bar{t}_{\mathsf{S}},\bar{t}_{\mathsf{D}}\right)=\begin{cases}Y & \text{if } X_{2,1}=Y \text{ and } y\geq\bar{t}_{\mathsf{S}}\\ Y & \text{if } X_{2,1}=Z \text{ and } y\geq m_2-\bar{t}_{\mathsf{D}}\\ Z & \text{otherwise.}\end{cases}$$

That is, agent 1 continues with his original technology Y (i) if both agents used the same technology and its value exceeds $\bar{t}_{\mathsf{S}}$; or (ii) if the agents used different technologies, but the other technology is either less valuable or its superior performance does not exceed by a margin larger than $\bar{t}_{\mathsf{D}}$ the value of the current technology. Let $\hat{\pi}_1\left(Y,X;\bar{t}_{\mathsf{S}},\bar{t}_{\mathsf{D}}\right)$ denote 1's reputation if he uses the double-threshold strategy, and adopts $X_{1,2}=X$ in period 2.

To see that an agent wants to distort the decision on $X_{1,2}$, suppose 1 were to use the first-best threshold values, $(\bar{t}_{\mathsf{S}},\bar{t}_{\mathsf{D}})=\left(\bar{y}_{\mathsf{dl}}^{FB},0\right)$. If 1 continues with his initial technology, the market deduces that either the same technology was used at the other site and its observed value exceeded $\bar{y}_{\mathsf{dl}}^{FB}$, or that the other site used the other technology which proved to be of inferior quality. Either event strengthens 1's reputation. Analogously, discontinuing a technology hurts a reputation. As a result, reputational concerns induce an agent to distort the decision in $t=2$. If both agents adopted Y in $t=1$, then agent 1 sticks to this technology if and only if $y+\lambda\hat{\pi}_1\left(Y,Y;\bar{t}_{\mathsf{S}},\bar{t}_{\mathsf{D}}\right)\geq E\left[\widetilde{Z}|s^Y,s^Y,y\right]+\lambda\hat{\pi}_1\left(Y,Z;\bar{t}_{\mathsf{S}},\bar{t}_{\mathsf{D}}\right)$. Similarly, in case agents adopted different technologies, agent 1 wants to continue with Y iff $y+\lambda\hat{\pi}_1\left(Y,Y;\bar{t}_{\mathsf{S}},\bar{t}_{\mathsf{D}}\right)\geq z+\lambda\hat{\pi}_1\left(Y,Z;\bar{t}_{\mathsf{S}},\bar{t}_{\mathsf{D}}\right)$. Proposition 2 describes equilibrium behaviour. Note that lo stands for locally determined reputations.

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 locally determined reputations, in equilibrium*
(i) truthful revelation is the communication strategy;
(ii) the belief functions are $\Pr\left(x_{2,1}|m_2\right)=1$ *for* $x_{2,1}=m_2$ *and* $\Pr\left(x_{2,1}|m_2\right)=0$ *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* $(\bar{t}_{\mathsf{S}}^*,\bar{t}_{\mathsf{D}}^*)$ *satisfy*

$$\lambda\left[\hat{\pi}_1\left(Y,Y;\bar{t}_{\mathsf{S}}^*,\bar{t}_{\mathsf{D}}^*\right)-\hat{\pi}_1\left(Y,Z;\bar{t}_{\mathsf{S}}^*,\bar{t}_{\mathsf{D}}^*\right)\right]=E\left[\widetilde{Z}|s^Y,s^Y,\bar{t}_{\mathsf{S}}^*\right]-\bar{t}_{\mathsf{S}}^* \quad (4)$$

$$\lambda\left[\hat{\pi}_1\left(Y,Y;\bar{t}_{\mathsf{S}}^*,\bar{t}_{\mathsf{D}}^*\right)-\hat{\pi}_1\left(Y,Z;\bar{t}_{\mathsf{S}}^*,\bar{t}_{\mathsf{D}}^*\right)\right]=\bar{t}_{\mathsf{D}}^*, \quad (5)$$

with $\bar{t}_{\mathsf{S}}^*\in\left(0,\bar{y}_{\mathsf{S}}^{FB}\right)$ *and* $\bar{t}_{\mathsf{D}}^*\in(0,1)$*. For* $\lambda\in[\underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}},\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}})$*, threshold values are* $(0,\bar{t}_{\mathsf{D}}^*)$ *and* $\bar{t}_{\mathsf{D}}^*$

solves $\lambda\hat{\pi}_1\left(Y,Y;0,\bar{t}^*_{\mathsf{D}}\right)=\bar{t}^*_{\mathsf{D}}$. *Finally, for* $\lambda\geq\bar{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}$, *threshold values equal* $(0,1)$.

Figure 2, panels (a) and (b) show the structure of the equilibrium. For $\lambda<\underline{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}$, see panel (a) and Eqs (4) and (5), at the equilibrium threshold values the size of the distortions, $E\left[\widetilde{Z}|s^Y,s^Y,\bar{t}^*_{\mathsf{S}}\right]-\bar{t}^*_{\mathsf{S}}$ and $\bar{t}^*_{\mathsf{D}}$, and the value of the reputational gap, $\lambda\left[\hat{\pi}_1\left(Y,Y;\bar{t}^*_{\mathsf{S}},\bar{t}^*_{\mathsf{D}}\right)-\hat{\pi}_1\left(Y,Z;\bar{t}^*_{\mathsf{S}},\bar{t}^*_{\mathsf{D}}\right)\right]$, are the same. The loss in technological value due to the distortion should in either case be compensated by the same boost in reputation. At $\lambda=\underline{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}$, $\bar{t}^*_{\mathsf{S}}=0$, and $\bar{t}^*_{\mathsf{D}}=E\left[\widetilde{Z}\right]$. The market infers from observing (Y,Z) that agents initially used different technologies and $y<z$, and so 1 initially picked the inferior technology, $\hat{\pi}_1\left(Y,Z;0,E\left[\widetilde{Z}\right]\right)=0$. Also, $\hat{\pi}_1\left(Y,Y;0,E\left[\widetilde{Z}\right]\right)>\pi$ as the market infers from (Y,Y) that either both agents initially received s^Y, or that the other agent received s^Z but $y\geq z-\bar{t}^*_{\mathsf{D}}$. Either possibility boosts agent 1's reputation. It follows from (4) that $\underline{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}=E\left[\widetilde{Z}\right]/\hat{\pi}_1\left(Y,Y;0,E\left[\widetilde{Z}\right]\right)<E\left[\widetilde{Z}\right]/\pi$.

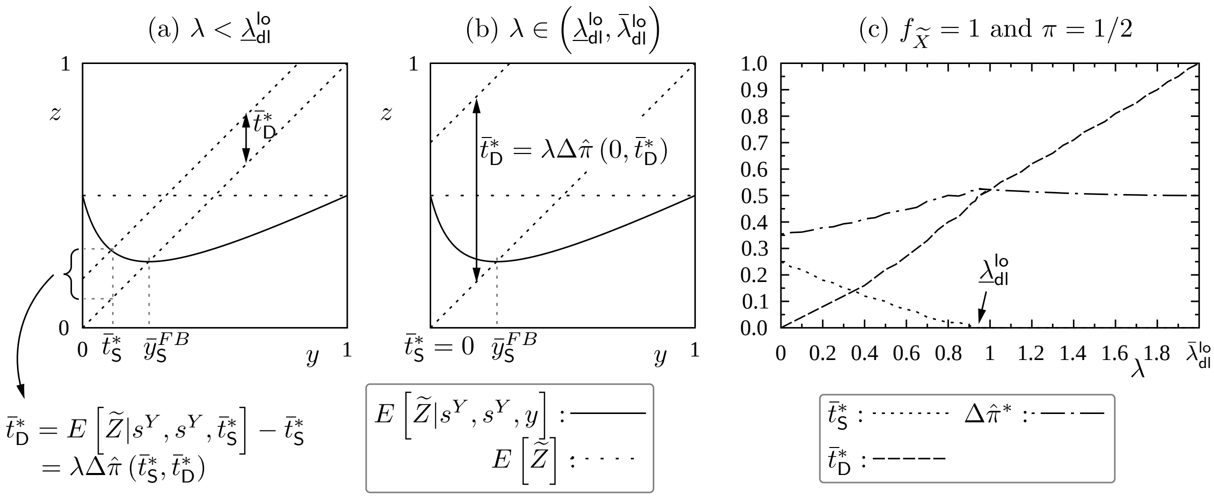

Figure 2: Decentralized decision making and locally determined reputations. Panels (a) and (b) depict the structure of equilibrium. Panel (c) reports equilibrium threshold values and the reputational gap for the uniform distribution and $\pi=1/2$. Hence, $\underline{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}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

For $\lambda\in[\underline{\lambda}^{\mathsf{lo}}_{\mathsf{dl}},\bar{\lambda}^{\mathsf{lo}}_{\mathsf{dl}})$, illustrated in panel (b), if 1 learns that 2 used the *same* technology, he continues his initial technology irrespective of its value y, $\bar{t}^*_{\mathsf{S}}=0$, whereas if 1 learns that 2 used a *different* technology, 1 may still change technology. For $\lambda\geq\bar{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}$, 1 sticks to his initial technology Y, irrespective of its value y, and regardless of what 2 reports, $(\bar{t}^*_{\mathsf{S}},\bar{t}^*_{\mathsf{D}})=(0,1)$. Then $\hat{\pi}_1\left(Y,Y;0,1\right)=\pi$ 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 $\pi=1/2$. The reputational gap rises for $\lambda<\underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$ to $\hat{\pi}_1\left(Y,Y;0,E\left[\widetilde{Z}\right]\right)>\pi$, and declines to π for $\underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}<\lambda\leq\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$.

5.2 Globally determined reputations

We start by showing that first-best behaviour, described on page 15, is not equilibrium behaviour. Suppose imputed equilibrium behaviour is first-best behaviour. Then, in $t=2$, either both agents adopt Y or both agents adopt Z. In the former case, the market infers that in $t=1$ agent 1 adopted the better technology and agent 2 the worse one. In the latter case, the inference is the reverse. Hence[17], $\hat{\pi}_1(Y,Z,Y,Y)=\frac{2\pi}{1+\pi}>\pi$ and $\hat{\pi}_1(Y,Z,Z,Z)=0$. Does an agent have an incentive to deviate from first-best behaviour? A unilateral deviation in the decision stage leads to a situation where the market observes that the agents use different technologies in $t=2$. The reputations implied by such a situation are 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 1 *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 $X_{2,2}=Y$).*

Now consider agent 1 in the decision stage of the game. Suppose that $y=m_2-\delta$, with $0<\delta<\lambda\hat{\pi}_1(Y,Z,Y,Z)$. Then, in the imputed equilibrium, $X_{1,2}=Z$, and agent 1's payoff equals m_2. Deviating and playing Y would yield $y+\lambda\hat{\pi}_1(Y,Z,Y,Z)>m_2$. That is, first-best behaviour is not equilibrium behaviour. We have proven the next Lemma.

Lemma 2 *First-best behaviour is not equilibrium behaviour in case of decentralized decision making with globally determined reputations.*

Having established that first-best behaviour is not equilibrium behaviour, we proceed with determining behaviour that does occur in equilibrium. Assume first that agents used

[17] See the proof of Proposition 3.

different technologies in $t = 1$, and consider the decision-making stage. With messages sent and beliefs set by the market, agent 1 continues with Y if and only if

$$
\begin{aligned}
& y + \lambda \Pr\left(X_{2,2} = Y | s^Y, s^Z, m_1, m_2\right) \hat{\pi}_1\left(Y, Z, Y, Y\right) \\
& +\lambda \Pr\left(X_{2,2} = Z | s^Y, s^Z, m_1, m_2\right) \hat{\pi}_1\left(Y, Z, Y, Z\right) \\
\geq \; & E\left[\widetilde{Z} | s^Y, s^Z, y, m_2\right] + \lambda \Pr\left(X_{2,2} = Y | s^Y, s^Z, m_1, m_2\right) \hat{\pi}_1\left(Y, Z, Z, Z\right) \\
& +\lambda \Pr\left(X_{2,2} = Z | s^Y, s^Z, m_1, m_2\right) \hat{\pi}_1\left(Y, Z, Z, Y\right).
\end{aligned}
\tag{6}
$$

As the conditional expected value $E\left[\widetilde{Z} | s^Y, s^Z, y, m_2\right] \equiv E\left[\widetilde{Z} | m_2\right]$ is independent of y, it follows that agent 1's decision strategy is a threshold strategy irrespective of the decision strategy of agent 2.[18] 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 agent 1 keeps Y and agent 2 switches to Y, then $y > z$. The implication is that agent 1 received a correct signal while agent 2 received an incorrect signal. Bayes' rule implies $\hat{\pi}_1\left(Y, Z, Y, Y\right) = \frac{2\pi}{1+\pi}$ and $\hat{\pi}_2\left(Y, Z, Y, Y\right) = 0$. Analogously, $\hat{\pi}_1\left(Y, Z, Z, Z\right) = 0$ and $\hat{\pi}_2\left(Y, Z, Z, Z\right) = \frac{2\pi}{1+\pi}$. 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. Application of Bayes' rule yields $\hat{\pi}_1\left(Y, Z, Y, Z\right) = \hat{\pi}_1\left(Y, Z, Z, Y\right) = \frac{\pi}{1+\pi}$. Using the posteriors and (6), agent 1 continues with Y if and only if

$$
y \geq E\left[\widetilde{Z} | m_2\right] - \lambda \frac{\pi}{1+\pi}, \tag{7}
$$

and analogously for agent 2. Note that if $m_2\left(z\right) = z$ for all z and $\lambda = 0$, the threshold strategy would coincide with first-best behaviour. Reputational concerns lead to a distortion in the technology decision in $t = 1$.

[18] That $E\left[\widetilde{Z} | s^Y, s^Z, y, m_2\right] \equiv E\left[\widetilde{Z} | m_2\right]$ is formally shown in Lemma A1 in the Appendix. Intuitively, as both agents are a priori equally likely to be of high ability, the conflicting signals cancel each other out and imply that there is no information on Z in the observed value of y over and above the information that is a priori present in y. With the prior distributions of Y and Z independent, the result follows.

Now consider the communication stage. In equilibrium, both agents anticipate that in the decision stage decisions are made according to (7). In particular, 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.[19]

In case agents initially adopted the *same* technology, Y, it is easy to see that truthful revelation is an equilibrium strategy. Communication is also irrelevant.[20] Proposition 3 below establishes that in this case an agent wants to deviate from first-best behaviour in the *decision* stage.

As communication breaks down in case of initial technologies differ, and is irrelevant in case of the same initial technology, the equilibrium decision strategy of 1 amounts to a comparison of y with a cut-off value that depends on the number of agents that used the same technology in $t = 1$. Let a *double-cut-off strategy* with cut-offs $(\bar{c}_{\mathsf{S}}, \bar{c}_{\mathsf{D}}) \geqq 0$ be defined as

$$d_1^{\mathsf{dl}}\left(I_1^{\mathsf{dl}}; \bar{c}_{\mathsf{S}}, \bar{c}_{\mathsf{D}}\right) = \begin{cases} Y & \text{if } X_{2,1} = Y \text{ and } y \geq \bar{c}_{\mathsf{S}} \\ Y & \text{if } X_{2,1} = Z \text{ and } y \geq \bar{c}_{\mathsf{D}} \\ Z & \text{otherwise.} \end{cases}$$

Of course, conditional on the information exchanged, the values of $\bar{c}_{\mathsf{S}}$ and $\bar{c}_{\mathsf{D}}$ that would maximize the technological value are $\bar{c}_{\mathsf{S}} = \bar{y}_{\mathsf{S}}^{FB}$, and $\bar{c}_{\mathsf{D}} = E\left[\widetilde{Z}\right]$. 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 globally determined reputations, there exists a unique equilibrium in which*

[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 $\hat{\pi}_1(Y, Y, Z, Z)$.

(i) the communication strategy is (a) a pooling strategy if initial technologies differ, and (b) truthful revelation if initial technologies are the same;
(ii) the belief function equals the density $f_1\left(z|I_1^{\mathsf{dl}}\right) = f(z)$ 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, $\bar{c}_{\mathsf{S}}^*$, satisfies

$$\lambda\left[\hat{\pi}_1\left(Y,Y,Y,Y;\bar{c}_{\mathsf{S}}^*\right) - \hat{\pi}_1\left(Y,Y,Z,Y\right)\right] = E\left[\widetilde{Z}|s^Y,s^Y,\bar{c}_{\mathsf{S}}^*\right] - \bar{c}_{\mathsf{S}}^*, \tag{8}$$

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

$$\lambda\frac{\pi}{1+\pi} = E\left[\widetilde{Z}\right] - \bar{c}_{\mathsf{D}}^*, \tag{9}$$

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

Figure 3, panels (a) and (b) correspond to (8) and (9), respectively.

Figure 3: Decentralized decision making and globally determined reputations. 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 $\hat{\pi}_1\left(Y,Y,Y,Y;\bar{c}_S^*\right) - \hat{\pi}_1\left(Y,Y,Z,Y\right)$, and $\Delta\hat{\pi}_1(\mathsf{D}) = \frac{\pi}{1+\pi}$.

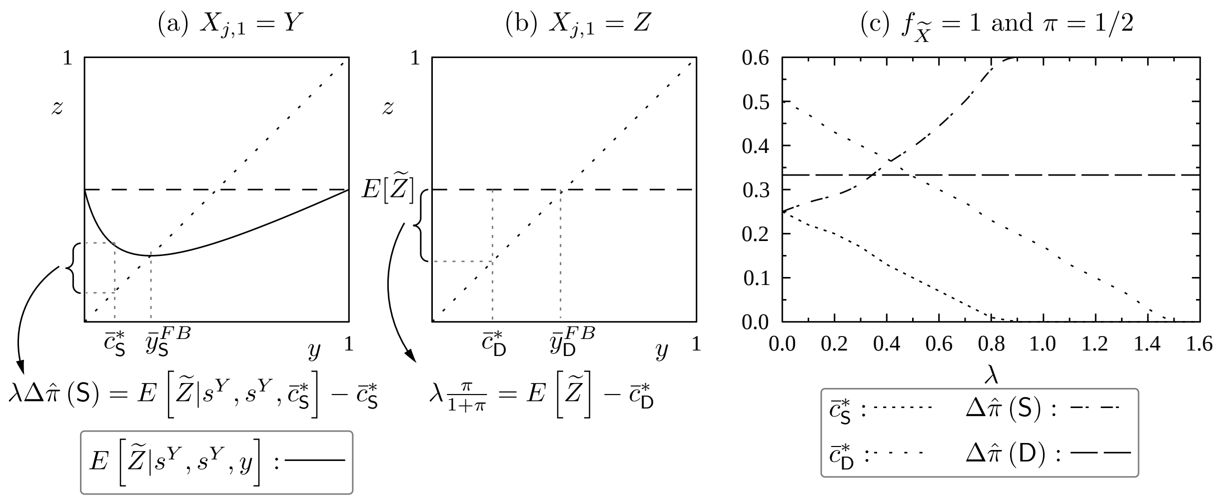

Panel (c) shows the equilibrium values in case of $f_X = 1$ and $\pi = 1/2$. Eq (9) shows that if agents adopted different technologies in $t = 1$, then the reputational gap is a constant

[22] The remark made in footnote 15 concerning $\bar{y}_{\mathsf{ia}}$ applies here to $\bar{c}_{\mathsf{S}}^*$.

function of $\bar{c}^*_{\mathsf{D}}$. 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 $\bar{c}^*_{\mathsf{D}}$ does not provide additional information on an agent's ability. If instead the market observes that agents initially adopted the *same* technology, $\hat{\pi}_1\left(Y, Y, Y, Y; \bar{c}^*_{\mathsf{S}}\right)$ does depend on the cut-off value: the lower is $\bar{c}^*_{\mathsf{S}}$, the lower is the reputation an agent commands in case of continuation.

In case reputations are globally determined, there are two reputational gaps. In panel (c), $\Delta\hat{\pi}_1\left(\mathsf{S}\right)$ is the reputational gap in case both agents started with Y. Starting with the same technology is a good sign for each agent's ability. Thus, even if λ increases and $\bar{c}^*_{\mathsf{S}}$ goes to zero, $\hat{\pi}_1\left(Y, Y, Y, Y; \bar{c}^*_{\mathsf{S}}\right) > \pi$. On the other hand, the lower is $\bar{c}^*_{\mathsf{S}}$, the more switching indicates a poor choice in period 1.

6 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,

$$d_C\left(I_C\right) = \begin{cases} Y, Y & \text{if } E\left[\widetilde{Y}|I_C\right] > E\left[\widetilde{Z}|I_C\right] \\ Y, Y & \text{if } E\left[\widetilde{Y}|I_C\right] = E\left[\widetilde{Z}|I_C\right] \text{ and } \mathsf{coin} = Y \\ Z, Z & \text{otherwise} \end{cases} \tag{10}$$

where "$\mathsf{coin}= Y$" means that the centre flips a fair coin with faces Y and Z, and Y comes up.[23] First best behaviour for each agent is to truthfully reveal his private information. We

[23] In a companion paper we analyse the case where both sites can continue with their initial technologies if $X_{1,1} = Y$ and $X_{2,1} = Z$ and $E\left[\widetilde{Y}|I_C\right] = E\left[\widetilde{Z}|I_C\right]$.

start by showing that first-best behaviour is not equilibrium behaviour in case of centralized learning. Each agent has an incentive to exaggerate the value of his technology.

Suppose imputed equilibrium behaviour is first-best behaviour. This requires that each agent truthfully reports the value of his technology to the centre. Truthfully reporting private information and (10) together imply that an agent commands a higher reputation if he is allowed to continue with "his" technology than if he is forced to change. That is, if reputations are determined locally, $\hat{\pi}_1(Y, Y) > \hat{\pi}_1(Y, Z)$, while if they are determined globally $\hat{\pi}_1(Y, Z, Y, Y) > \hat{\pi}_1(Y, Z, Z, Z)$. We now show that agent 1 has an incentive to slightly exaggerate the value of his technology: for all parameter values, there is an $\varepsilon > 0$ such that sending $m_1(y) = y + \varepsilon$ yields a higher expected payoff than sending $m_1 = y$. To see this, suppose (10), the above posteriors and that agent 2 truthfully reports. For $z \notin (y, y + \varepsilon)$, sending $m_1 = y + \varepsilon$ rather than $m_1 = y$ does not affect agent 1's payoff. For $z \in (y, y + \varepsilon)$, the reputational benefit of sending $m_1 = y + \varepsilon$ rather than $m_1 = y$ is positive ($\lambda[\hat{\pi}_1(Y, Y) - \hat{\pi}_1(Y, Z)] > 0$ for local reputations and $\lambda[\hat{\pi}_1(Y, Z, Y, Y) - \hat{\pi}_1(Y, Z, Z, Z)] > 0$ for globally determined reputations) and independent of ε. The loss of using the inferior technology due to sending $m = y + \varepsilon$ can be made arbitrarily small by reducing the value of ε. This shows that a profitable deviation from first-best behaviour exists.

Lemma 3 *Under centralized decision making, an equilibrium in which agents truthfully reveal their private information does not exist, neither in case of locally nor in case of globally determined reputations.*

We show that in case agents start out with different technologies all 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 intervals, 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.

Interestingly, we also show that in case of globally determined reputations and for any finite weight λ that the agent attaches to his reputation, influential communication remains possible, and that a babbling equilibrium is not neologism-proof.

Formally, let $\mathbf{a}(N) \equiv (a_0(N), \ldots, a_N(N))$ denote a partition of $[0,1]$ in N intervals, with $0 = a_0(N) < a_1(N) < \cdots < a_N(N) = 1$. Agent 1 is said to use a *partition strategy* to communicate if there exists a tuple $(N, \mathbf{a}(N))$, such that $\mu_1^{\mathsf{P}}\left(m_1 | s^Y, y, X_{2,1}\right)$ is uniform, supported on $[a_r(N), a_{r+1}(N)]$ if $y \in (a_r(N), a_{r+1}(N))$ for $r = 0, \ldots, N-1$.[24] The contents of these messages – what they imply concerning the expected value of the technology – are the same if $m_1, m_2 \in [a_{r-1}, a_r)$ and they differ if $m_1 < a_r \leq m_2$ for some r. We focus on the highest value of N consistent with incentives. 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, $N \geq 2$. To save space, we write $\mathbf{a}$ instead of $\mathbf{a}(N)$ if this does not lead to confusion.

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 interest are different from those of the centre, but identical to those of the 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 such that the centre will next decide that the technology is sufficiently good to merit continuation, and one 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.[25] 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, $\bar{y}_{\mathsf{S}} \in [0,1]$, for which an agent is indifferent between sending one message rather than the other.

The next proposition characterizes equilibrium behaviour. To state the belief function of the centre, define a truncated density as follows: $Tr(x; a_r, a_{r+1}) = g(x) / (F(a_{r+1}) - F(a_r))$,

[24] Note that in any interval of the partition the expert uses a random strategy. This guarantees that in equilibrium any possible message is sent with strictly positive probability. A discussion of out-of-equilibrium beliefs (what should the planner think about the value of a technology if he were to observe a non-equilibrium message?) can thus be avoided.

[25] In a previous version of this paper we show that truthfully revealing information to the center in case agents use the same technology can be part of equilibrium. However, it amounts to playing a weakly dominated strategy, an unlikely candidate to describe agents' behaviour.

where $g(x) = f(x)$ for $x \in [a_r, a_{r+1}]$ and $g(x) = 0$ everywhere else.

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

(i) the centre's decision strategy is as defined in (10).

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

(ii) the centre's belief function is (a) $f_C(y|I_C) = Tr\left(y; a_r^*, a_{r+1}^*\right)$ *for* $m_1 \in \left(a_r^*, a_{r+1}^*\right)$ *and* $f_C(z|I_C) = Tr\left(z; a_r^*, a_{r+1}^*\right)$ *for* $m_2 \in \left(a_r^*, a_{r+1}^*\right)$, *for* $r = 0, \ldots, N^* - 1$ *if initial technologies differ, and (b)* $f_C(y|I_C) = Tr(y; 0, \bar{y}_{\mathsf{S}}^*)$ *for* $m_1 \in [0, \bar{y}_{\mathsf{S}}^*]$ *and* $f_C(x_{1,1}|I_C) = Tr(y; \bar{y}_{\mathsf{S}}^*, 1)$ *for* $m_1 \in (\bar{y}_{\mathsf{S}}^*, 1]$ *if initial technologies are the same;*

(iii) in case of locally determined reputations, *the partition* $\mathbf{a}^*$ *and the collusion strategy* $\bar{y}_{\mathsf{S}}^* = \bar{y}_{\mathsf{S}}^{\mathsf{lo}*}$ *satisfy*

$$\begin{aligned}
\lambda\left[\hat{\pi}\left(Y, Y; \bar{y}_{\mathsf{S}}^{\mathsf{lo}*}, \mathbf{a}^*\right) - \hat{\pi}\left(Y, Z; \bar{y}_{\mathsf{S}}^{\mathsf{lo}*}, \mathbf{a}^*\right)\right] &= E\left[\widetilde{Z}|a_{r-1}^* \leq z \leq a_{r+1}^*\right] - a_r^* \qquad (11) \\
\lambda\left[\hat{\pi}\left(Y, Y; \bar{y}_{\mathsf{S}}^{\mathsf{lo}*}, \mathbf{a}^*\right) - \hat{\pi}\left(Y, Z; \bar{y}_{\mathsf{S}}^{\mathsf{lo}*}, \mathbf{a}^*\right)\right] &= E\left[\widetilde{Z}|s^Y, s^Y, \bar{y}_{\mathsf{S}}^{\mathsf{lo}*}\right] - \bar{y}_{\mathsf{S}}^{\mathsf{lo}*} \qquad (12)
\end{aligned}$$

for $r = 1, \ldots, N^* - 1$. *For* $\lambda < \bar{\lambda}_{\mathsf{cl}}^{\mathsf{lo}}$, $N^* \geq 2$ *and* $\bar{y}_{\mathsf{S}}^{\mathsf{lo}*} > 0$, *whereas for* $\lambda \geq \bar{\lambda}_{\mathsf{cl}}^{\mathsf{lo}}$, $N^* = 1$ *and* $\bar{y}_{\mathsf{S}}^{\mathsf{lo}*} = 0$. *That is, the agents do not send influential information on* y *and* z *for* $\lambda \geq \bar{\lambda}_{\mathsf{cl}}^{\mathsf{lo}}$.

(iv) in case of globally determined reputations, *the partition* $\mathbf{a}^*$ *satisfies*

$$\lambda\left[\hat{\pi}_1(Y, Z, Y, Y; \mathbf{a}^*) - \hat{\pi}_1(Y, Z, Z, Z; \mathbf{a}^*)\right] = E\left[\widetilde{Z}|a_{r-1}^* \leq z \leq a_{r+1}^*\right] - a_r^* \qquad (13)$$

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

$$\lambda\left[\hat{\pi}\left(Y, Y, Y, Y; \bar{y}_{\mathsf{S}}^{\mathsf{gl}*}\right) - \hat{\pi}\left(Y, Y, Z, Z; \bar{y}_{\mathsf{S}}^{\mathsf{gl}*}\right)\right] = E\left[\widetilde{Z}|s^Y, s^Y, \bar{y}_{\mathsf{S}}^{\mathsf{gl}*}\right] - \bar{y}_{\mathsf{S}}^{\mathsf{gl}*}. \qquad (14)$$

Moreover, for any finite λ, $N^* \geq 2$. *For* $\lambda < \bar{\lambda}_{\mathsf{cl}}^{\mathsf{gl}}$ $\bar{y}_{\mathsf{S}}^{\mathsf{gl}*} > 0$, *whereas for* $\lambda \geq \bar{\lambda}_{\mathsf{cl}}^{\mathsf{gl}}$ $\bar{y}_{\mathsf{S}}^{\mathsf{gl}*} = 0$. *That is, in case agents initially used different strategies, agents send influential information about* y *and* z *for any finite* λ. *Moreover, a complete pooling (babbling) equilibrium is not neologism-proof. If agents initially used the same technology, they do not send influential information about the technology's values for* $\lambda \geq \bar{\lambda}_{\mathsf{cl}}^{\mathsf{gl}}$.

In case agents used different technologies, the communication strategy is a partition strategy. The idea behind the partition strategy is as follows. After an agent has observed the value of his technology, he decides how to "rank" it. The higher the rank is, the more likely it becomes that the centre chooses his technology. This suggests that his technology is the better one. As a result, agent 1 enjoys a reputational benefit. Ranking the technology high also has a cost. If $z > y$ but agent 2 does not rank Z as high as 1 ranks Y, the centre forces both agents to choose Y, the inferior technology, in period 2. This possibility stops the agent from ranking his technology too high. The left-hand sides of the equations (11) and (13) state the net reputational value of continuing with one's technology. For $y = a_r^*$, this gain is exactly offset by a loss in expected project value due to continuation: the agent is indifferent between using two adjacent ranks (messages) to describe the value of technology Y. Sending one message rather than the other changes the choice of the centre only for $z \in \left(a_{r-1}^*, a_{r+1}^*\right)$, see the right-hand side of (11) and (13).

In terms of informativeness, a partition strategy is in between the truthful revelation that characterizes communication in case of decentralized decision making with *locally* determined reputations and the absence of communication in case of decentralized decision making process and *globally* determined reputations. 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 globally determined reputations, the loss of decision-making power makes that an agent becomes cautious when communicating: his own exaggerated claims are no longer costless but can lead to an inferior choice at the agent's own site.

The striking result in case of globally determined reputations 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 of cardinality two. This partition is fully determined by $a_1(2) = a_1$, and Eq (13) reduces to[26]

$$\lambda \frac{4\pi}{1+\pi} F(a_1)\left(1 - F(a_1)\right) = E\left[\widetilde{Z}\right] - a_1. \tag{15}$$

[26] See the proof of Proposition 4.

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 globally determined reputations if agents start out with different technologies.

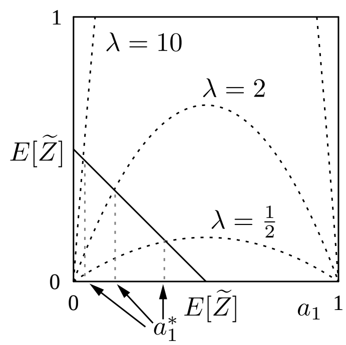

The dotted lines represent the left-hand side of Eq (15) for various values of λ. The drawn line graphs the right-hand side. Clearly, for any finite λ, $a_1^* > 0$ such that influential communication is part of equilibrium behaviour.

To understand why for $a_1 = 0$ the reputional 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 *locally* determined reputations, 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 $\bar{y}_{\mathsf{S}}^{\mathsf{lo}*} = 0$ the reputational gap does *not* vanish but equals $\frac{4\pi^2}{(3+\pi^2)(1+\pi)}$.[27] 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 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 $a_1 = 0$.

[27] See the proof of Proposition 4.

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, $E\left[\widetilde{Y}|I_C^{\mathsf{cl}}\right] = E\left[\widetilde{Y}\right]$ and $E\left[\widetilde{Z}|I_C^{\mathsf{cl}}\right] = E\left[\widetilde{Z}\right]$. Proposition 4 shows that 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]

7 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 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

[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.

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,

$$W(\lambda, \pi) = \frac{E\left[\widetilde{X}_{1,2}|s^Y, \lambda, \pi\right] - \left(\pi E\left[\widetilde{Y}|y > z\right] + (1-\pi) E\left[\widetilde{Y}\right]\right)}{E\left[\widetilde{Y}|y > z\right] - \left(\pi E\left[\widetilde{Y}|y > z\right] + (1-\pi) E\left[\widetilde{Y}\right]\right)} * 100\%. \tag{16}$$

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.'

7.1 Decentralized decision making: welfare comparisons

A key finding of this subsection is that *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 in case of decentralized decision making and globally determined reputations agent 1 also knows $X_{2,1}$ (but not $x_{2,1}$ if $X_{2,1} = Z$), and that if reputations are locally determined he knows both $X_{2,1}$ 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, $W_{\mathsf{ia}}(\lambda, \pi) < W_{\mathsf{dl}}^{\mathsf{gl}}(\lambda, \pi) < W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda, \pi)$.[29] However, propositions 2 and 3 imply that 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 $\bar{\lambda}_{\mathsf{ia}} = E\left[\widetilde{Z}\right]/\pi$, $\bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}} = 1/\pi$, and

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

$\bar{\lambda}^{\mathsf{gl}}_{\mathsf{dl}} = E\left[\widetilde{Z}\right](1+\pi)/\pi$ for an isolated agent, an agent whose reputation is locally determined, and an agent whose reputation is globally determined, respectively. As $\bar{\lambda}_{\mathsf{ia}} < \bar{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}, \bar{\lambda}^{\mathsf{gl}}_{\mathsf{dl}}$ for all parameter values, an isolated agent stops using information for a lower value of λ than a decentralized agent.[30] Interestingly, for $E\left[\widetilde{Z}\right](1+\pi) > 1$, $\bar{\lambda}^{\mathsf{gl}}_{\mathsf{dl}} > \bar{\lambda}^{\mathsf{lo}}_{\mathsf{dl}}$ holds. In other words, for sufficiently high values of λ, if markets observe the choices both agents make, then, compared with a situation in which markets only know the actions of a single agent, welfare goes up, although agents have communicated less information among each other. Less information for decision-makers may mean higher welfare.

The reason is that information has two roles. On the one hand, additional information helps the *agent* in identifying the better technology; on the other, additional information helps the *market* in evaluating an agent's ability. Consider the first role. With locally determined reputations, agent 1 knows the value of the other technology. The difference $|z-y|$ can be as large as 1. In case of globally determined reputations, agent 1 does not know the value of the other technology. Instead, he can only calculate $\left|E\left[\widetilde{Z}\right]-y\right|$. This difference is strictly less than one. Hence, ceteris paribus, for all information about Z to be ignored and for the agent to continue with Y, λ should be larger when his reputation is locally determined than when it is globally determined. Now turn to the second role. If a market cannot compare across sites, then reputation-wise more is at stake when the agent takes a decision relative to the case where a market can compare across sites. In the latter case the market already knows whether agents used the same or different technologies. If agent 1 were to continue with Y, rather than to switch to Z, independent of what he knows about Z, then the reputational gap equals π if markets have access to local information only and $\pi/(1+\pi) < \pi$ if markets can compare across sites. Hence, ceteris paribus, λ should be larger in the latter case than in the former case for information about Z to be ignored and for the agent to continue with Y. The inequality $E\left[\widetilde{Z}\right](1+\pi) > 1$ 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.

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

In Figure 5, we depict W, for decentralized decision making with reputations that are locally and globally determined and for isolated agents under the assumption that the value of technology $X \in \{Y, Z\}$ is uniformly distributed, $f_X(x) = 1$ on $[0,1]$, and that $\pi = \frac{1}{2}$. [31]

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

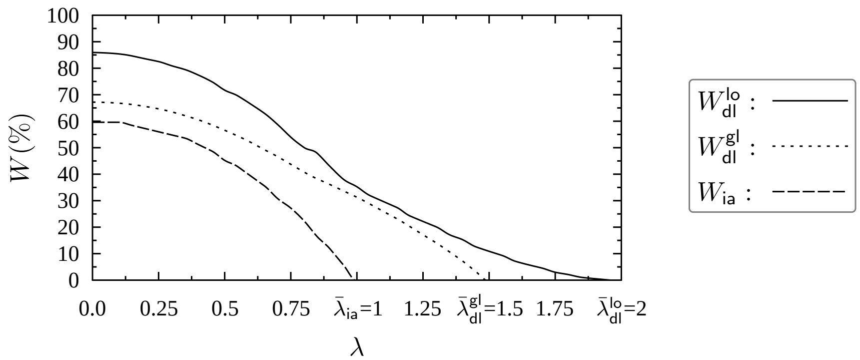

Figure 5 illustrates a number of points. First, learning from one's own past behaviour and from others potentially boosts welfare enormously. For λ close to zero, an isolated agent who is of high ability with probability $\pi = 1/2$ and learns from his own experience only can capture 60% of the increase in expected project value. Learning from others further increases this percentage. Second, 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. Third, the relative performance does not change in λ. Additional calculations (not reported here) show that this is true independent of the value of π. The following Proposition sums up.[32]

Proposition 5 *For any f_X and π, there exists a $\lambda_1 > 0$ such that $W_{\mathsf{ia}}(\lambda, \pi) < W^{\mathsf{gl}}_{\mathsf{dl}}(\lambda, \pi) < W^{\mathsf{lo}}_{\mathsf{dl}}(\lambda, \pi)$ 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 $W_{\mathsf{ia}}(\lambda, \pi) < W^{\mathsf{gl}}_{\mathsf{dl}}(\lambda, \pi) < W^{\mathsf{lo}}_{\mathsf{dl}}(\lambda, \pi)$ 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 $W_{\mathsf{ia}}(\lambda, \pi) <$*

[31] For $f_X = 1$ and $\pi = 1/2$, $E\left[\widetilde{Y}|y > z\right] = 2/3$ and $\pi E\left[\widetilde{Y}|y > z\right] + (1-\pi)E\left[\widetilde{Y}\right] = 7/12$.

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

$W^{\mathsf{lo}}_{\mathsf{dl}}(\lambda,\pi) < W^{\mathsf{gl}}_{\mathsf{dl}}(\lambda,\pi)$ *for* $\lambda > \lambda_3$. *For* $f_X = 1$, *the uniform distribution,* $W_{\mathsf{ia}}(\lambda,\pi) < W^{\mathsf{gl}}_{\mathsf{dl}}(\lambda,\pi) < W^{\mathsf{lo}}_{\mathsf{dl}}(\lambda,\pi)$ *holds for all* λ *and* π.

7.2 Centralized decision making: welfare comparisons

A key finding in this subsection is that for sufficiently large values of λ, welfare is unambiguously higher if the centre receives her information from agents whose reputations are globally determined. This stems from the fact, reported in Proposition 4, that if markets are aware of the technologies used by both agents, these agents send influential information for all values of the parameters, while they will babble for sufficiently high values of λ if comparisons across agents are not possible.

Propositions 1 and 4 allow us to compare welfare W in case of centralized decision making and isolated agents.[33]

Proposition 6 *For any* f_X *and* π, *there exists a* $\lambda_4 > 0$ *such that* $W_{\mathsf{ia}}(\lambda,\pi) < W^{\mathsf{lo}}_{\mathsf{cl}}(\lambda,\pi), W^{\mathsf{gl}}_{\mathsf{cl}}(\lambda,\pi)$ *for all* $\lambda \in (0,\lambda_4)$. *Furthermore, for any* f_X *and* π *there exists a* $\lambda_5 > 0$ *such that* $W_{\mathsf{ia}}(\lambda,\pi) < W^{\mathsf{lo}}_{\mathsf{cl}}(\lambda,\pi) < W^{\mathsf{gl}}_{\mathsf{cl}}(\lambda,\pi)$ *for all* $\lambda > \lambda_5$. *In addition, for* $f_X = 1$, *the uniform distribution,* $W_{\mathsf{ia}}(\lambda,\pi) < W^{\mathsf{gl}}_{\mathsf{cl}}(\lambda,\pi), W^{\mathsf{lo}}_{\mathsf{cl}}(\lambda,\pi)$ *holds for all* $\lambda > 0$ *and* π.

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. $f_{\widetilde{X}} = 1$ and $\pi = 1/2$, such that $\bar{\lambda}_{\mathsf{ia}} = 1$, $\bar{\lambda}^{\mathsf{lo}}_{\mathsf{cl}} = 2\frac{7}{16}$.

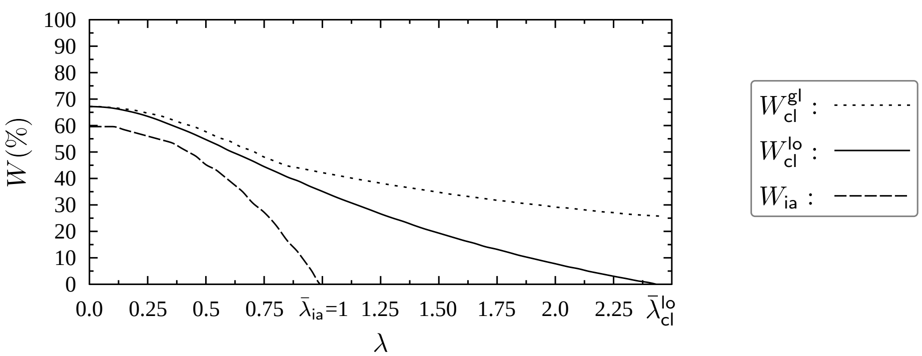

[33] If $\lambda \geq \bar{\lambda}^{\mathsf{lo}}_{\mathsf{cl}}$, then, $W_{\mathsf{ia}}(\lambda,\pi) = W^{\mathsf{lo}}_{\mathsf{cl}}(\lambda,\pi) = 0$. This case is ignored in Proposition 6.

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.

7.3 Further comparisons

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 turn to the complementary question, and analyse, for a given information base on which reputations can be based, the conditions that determine whether a decentralized process or a centralized process performs best.

If reputations are locally determined, the learning process that is best depends fundamentally on the parameters of the model.[34]

Proposition 7 *Suppose reputations are locally determined. Whether centralized or decentralized decision making creates more welfare for high values of λ depends on the sign of*

$$\frac{1}{E\left[\widetilde{Z}\right]} - \frac{\left(3+\pi^2\right)\left(1+\pi\right)}{4\pi}. \tag{17}$$

If the sign is positive (negative), welfare is higher with decentralized (centralized) decision making for sufficiently high values of λ.

[34] If $\lambda \geq \bar{\lambda}_{\mathsf{cl}}^{\mathsf{lo}}$, then $W_{\mathsf{dl}}^{\mathsf{lo}}(\lambda, \pi) = W_{\mathsf{cl}}^{\mathsf{lo}}(\lambda, \pi)$. This case is ignored in Proposition 7.

If the sign of the expression in (17) is negative, there are values of λ such that under a decentralized process the technology adoption decision in $t = 2$ depends on the observed values y and z, whereas in a centralized process, agents do not transmit useful information about their technologies to the centre. As a result, expected welfare is higher in case of decentralized decision making.

Note that $(3 + \pi^2)(1 + \pi)/4\pi > 2$ for all π. $1/E\left[\widetilde{Z}\right]$ is the ratio of the upperbound of the support and the expected value of the technology. Hence, the ratio should exceed 2 for there to be values of π such that decentralization outperforms centralization for high values of λ. The uniform distribution cannot meet this condition. Does this mean that welfare is higher under a centralized process than under a decentralized process in case of the uniform distribution for *all* λ and π? The next proposition provides sufficient conditions on λ and π such that a decentralized process outperforms a centralized process.

Proposition 8 *Assume reputations are locally determined and technology values are uniformly distributed. Then, for any π there are values of λ, $\lambda \in \left[\underline{\lambda}(\pi), \bar{\lambda}(\pi)\right]$, with $0 < \underline{\lambda}(\pi) < \bar{\lambda}(\pi)$, such that welfare $W(\lambda, \pi)$ is higher under decentralized than under centralized decision making.*

What learning process is best if reputations can also be based on comparisons across sites?

Proposition 9 *In case of globally determined reputations, 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 markets observe technology adoption at both sites 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 a_1^* 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).

8 Discussion

8.1 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 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.[35] This variation has been a source of worries as, most likely, some patients do not receive optimal treatment.[36] 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.[37] 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.[38] It also seems that the IT revolution and increased information dissemination

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

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

[37] See Fink et al. (1984) and Rowe et al. (1991).

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

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 locally versus globally determined 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 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.[39] 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.[40] 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

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

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

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.

8.2 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 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 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.

8.3 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 locally determined reputations. 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.

History. Our approach is particularly relevant for situations where agents *independently* gain experience that is worth sharing. In our model, period 1 represents history. However, in other situations experience still has to be gained. Consider therefore a different institutional set up than the ones discussed above in which decision rights are centralized in period 1. This centre could then opt for ignoring signals and assign one technology to agent 1 and the other technology to agent 2. Such a procedure is likely to weaken reputational concerns as the technology decisions are no longer linked to signals. Moreover, it allows for learning in period 2. It is easy to show that assigning technologies in period 1 is optimal if signals are not very informative. The first-period costs of ignoring signals are then small.

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.

Communication in case of decentralized decision making and locally determined reputations. We have assumed throughout the paper that an agent cares about the value of a technology and his reputation. We noticed that in the case of locally determined reputations and decentralized decision making any communication strategy can be an equilibrium strategy. An infinitesimal aversion to lying helps to uniquely pin down the communication strategy. It would be interesting to test in the laboratory whether subjects would actually tell the truth or would use a different communication strategy. In an experiment in which subjects are medical specialists that choose patients' treatments, would a subject opt for informing a fellow specialist as accurately as possible or would he take into account the other specialist's tendency to hang on to his first-period intervention and thus exaggerate the value of the technology in the interests of the latter fellow's patients?

Appendix

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 $\Pr\left(\bar{\theta}|s^Y,y\right)=2F(y)\pi/\left(2F(y)\pi+(1-\pi)\right)$. 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. ■

Proof of Proposition 1: First, $\hat{\pi}(YY;\bar{t})=\Pr\left(\bar{\theta}|YY;\bar{t}\right)$ in (2). Use $\Pr\left(YY|\bar{\theta}\right)=\Pr\left(y\geq\bar{t}|\bar{\theta}\right)/2=\left(1-F(\bar{t})^2\right)/2$ and $\Pr\left(YY|\underline{\theta}\right)=\Pr\left(y\geq\bar{t}|\underline{\theta}\right)/2=\left(1-F(\bar{t})\right)/2$, and apply Bayes rule (analogously for $\hat{\pi}(YZ;\bar{t})$. 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 $\bar{t}=\bar{y}^*_{\mathsf{ia}}\leq\bar{y}^{FB}_{\mathsf{ia}}$. To see that $\bar{y}^*_{\mathsf{ia}}$ is a decreasing function of λ for $\lambda\leq\bar{\lambda}_{\mathsf{ia}}$, define $\delta:=\bar{y}^{FB}_{\mathsf{ia}}-\bar{y}_{\mathsf{ia}}$ 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 $\bar{t}_{\mathsf{S}}^* = 0$, the RHS of (4) equals $E\left[\widetilde{Z}\right]$, and therefore $\bar{t}_{\mathsf{D}}^* = E\left[\widetilde{Z}\right]$, and thus $\lambda = \underline{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$. Finally, if $\bar{t}_{\mathsf{S}}^* = 0$ 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 $\bar{t}_{\mathsf{D}}^* \uparrow 1$) such that for $\lambda \geq \bar{\lambda}_{\mathsf{dl}}^{\mathsf{lo}}$, the agents indeed continue with their initial technologies no matter what.■

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 $\Omega := \left\{s^Y, s^Z, \bar{y}, z \in A\right\}$, and write

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

Write $\Pr(\Omega|sm, db) = \Pr(\Omega|sm, db, y > z)\Pr(y > z) + \Pr(\Omega|sm, db, y < z)\Pr(y < z)$, 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

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

and analogously $\Pr(\Omega|sm,db,z>y) = \Pr(\Omega|db,sm,y>z) = 0$, $\Pr(\Omega|db,sm,z>y) = f(\bar{y})\int_{HA(\bar{y})} f(z)\,dz$, $\Pr(\Omega|db,db,y>z) = \frac{1}{2}f(\bar{y})\int_{LA(\bar{y})} f(z)\,dz$, and $\Pr(\Omega|db,db,z>y) = \frac{1}{2}f(\bar{y})\int_{HA(\bar{y})} f(z)\,dz$. Thus, $\Pr(\Omega|sm,db) = \frac{1}{2}f(\bar{y})\int_{LA(\bar{y})} f(z)\,dz$, $\Pr(\Omega|db,sm) = \frac{1}{2}f(\bar{y})\int_{HA(\bar{y})} f(z)\,dz$ and $\Pr(\Omega|db,db) = \frac{1}{4}f(\bar{y})\int_{A} f(z)\,dz$. Therefore, $\Pr(\Omega) = \frac{1}{4}(1-\pi^2)f(\bar{y})\int_A f(z)\,dz$. 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 $\{s^Y, s^Z, \bar{y}, z\in A, sm, db\}$ implies that $z<y$ and $z\in A$. We will write this event as "$\Omega; sm, db$". Thus,

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

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

Proof of Proposition 3: First, the reputations. $\hat{\pi}_1(YYYY;\bar{c}) = \Pr(\bar{\theta}|YYYY;\bar{c})$. Write $F(\bar{c}) = F$. Use $\Pr(YYYY|\bar{\theta}) = \Pr(YYYY|\bar{\theta}, y>z)/2 = (1+\pi)\Pr(y>\bar{c}|y>z)/4 = (1+\pi)(1-F^2)/4$ 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 $\bar{c}^*_{\mathsf{S}} = 0$, the gap equals $\frac{1+\pi}{1+\pi^2}\pi$. Also, $\hat{\pi}_1(YZYY) = \Pr(\bar{\theta}|YZYY)$. From $\{Y,Z,Y,Y\}$ the market deduces that $y>z$ in case of both first-best and equilibrium behaviour. Thus, $\Pr(YZYY|\bar{\theta}) = (1-\pi)/4$ (as $\theta_2 = \underline{\theta}$ for $X_{2,1} = Z$) and $\Pr(YZYY|\underline{\theta}) = (1-\pi)/8$, and apply Bayes rule. Finally, $\hat{\pi}_1(YZYZ) = \Pr(\bar{\theta}|YZYZ)$. From $\{Y,Z,Y,Z\}$ the market deduces that $(y,z)\in A := \{(y,z)\,|y,z<\bar{c}^*_{\mathsf{D}} \text{ or } y,z>\bar{c}^*_{\mathsf{D}}\}$. Use $\Pr(YZYZ|\bar{\theta}) = \frac{1-\pi}{2}\Pr(A|y>z)\frac{1}{2}$,

$\Pr(YZYZ|\underline{\theta}) = \frac{1}{2}\frac{1-\pi}{2}\Pr(A|y>z)\frac{1}{2} + \frac{1}{2}\frac{1+\pi}{2}\Pr(A|z>y)\frac{1}{2}$, and $\Pr(A|z>y) = \Pr(A|y>z)$ (as Y and Z are iid), and apply Bayes rule. For given reputations and behaviour of 2, if $y = \overline{c}^*_{\mathsf{D}}$, and if 1 continues Y he gets $\overline{c}^*_{\mathsf{D}} + \lambda \Pr(z < \overline{c}^*_{\mathsf{D}})\, 2\pi/(1+\pi) + \lambda \Pr(z \geq \overline{c}^*_{\mathsf{D}})\, \pi/(1+\pi)$, whereas switching to Z yields $E\left[\widetilde{Z}\right] + \lambda \Pr(z < \overline{c}^*_{\mathsf{D}})\, \pi/(1+\pi)$. Equating these expressions, one obtains (9). It is immediate that $\overline{c}^*_{\mathsf{D}}$ is a decreasing function of λ. The comparative statics result on $\overline{c}^*_{\mathsf{S}}$ uses Theorem 3 in Milgrom and Roberts (1994), see also proof of Proposition 1. The expressions for $\underline{\lambda}^{\mathsf{gl}}_{\mathsf{dl}}$ and $\overline{\lambda}^{\mathsf{gl}}_{\mathsf{dl}}$ are then immediate.■

Proof of Proposition 4: (A) Consider $X_{1,1} \neq X_{2,1}$. Assume that the centre uses (10). As $I^{\mathsf{cl}}_C = \left\{s^Y, s^Z, m^Y, m^Z\right\}$, $E\left[\tilde{Y}|I^{\mathsf{cl}}_C\right] = E\left[\tilde{Y}|m^Y\right]$ and $E\left[\tilde{Z}|I^{\mathsf{cl}}_C\right] = E\left[\tilde{Z}|m^Z\right]$. 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 utitiliy 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^Y_1 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 $H\left(m^Y\right) = \left\{z : E\left[\tilde{Y}|m^Y\right] > E\left[\tilde{Z}|m^Z(z)\right]\right\}$, $G\left(m^Y\right) = \left\{z : E\left[\tilde{Y}|m^Y\right] = E\left[\tilde{Z}|m^Z(z)\right]\right\}$, and $L\left(m^Y\right) = \left\{z : E\left[\tilde{Y}|m^Y\right] < E\left[\tilde{Z}|m^Z(z)\right]\right\}$, and the probabilities $p\left(m^Y\right) := \int_{H(m^Y)} f(z)\,dz$ and $q\left(m^Y\right) := \int_{G(m^Y)} f(z)\,dz$. Hence, if 1 sends message m_1^Y if he has observed y, then in equilibrium his expected payoffs equal

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

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 $\hat{\pi}\left(Y, X; \bar{y}_{\mathsf{S}}^{\mathsf{lo}*}, \mathbf{a}^*\right)$. 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

$$\begin{aligned} & p(a_{r-1}, a_r)[a_r + \lambda\hat{\pi}_1(YY)] + \frac{1}{2}p(a_r, a_{r+1})[a_r + \lambda\hat{\pi}_1(YY)] + \\ & \frac{1}{2}p(a_r, a_{r+1})\left[E\left[\widetilde{Z}|a_r \leq z < a_{r+1}\right] + \lambda\hat{\pi}_1(YZ)\right], \end{aligned} \tag{18}$$

whereas m_r yields

$$\begin{aligned} & \frac{1}{2}p(a_{r-1}, a_r)[a_r + \lambda\hat{\pi}_1(YY)] + \frac{1}{2}p(a_{r-1}, a_r)\left[E\left[\widetilde{Z}|a_{r-1} \leq z < a_r\right] + \lambda\hat{\pi}_1(YZ)\right] \\ & + p(a_r, a_{r+1})\left[E\left[\widetilde{Z}|a_r \leq z < a_{r+1}\right] + \lambda\hat{\pi}_1(YZ)\right]. \end{aligned} \tag{19}$$

Equating (18) and (19) shows that agent 1 is indifferent between sending m_{r+1} and m_r for $y = a_r$ if (11) holds.

(B) If $X_{1,1} = X_{2,1} = Y$, 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 $\bar{y}^{\mathsf{lo}*}_{\mathsf{S}}$ satisfying (12) 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 $\bar{y}^{\mathsf{lo}*}_{S} = 0, \mathbf{a}^*_1(1) = 0$ (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 (11) and (12) 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, $\hat{\pi}\left(YYYY; \bar{y}^{\mathsf{gl}*}_{\mathsf{S}} = 0\right) = (1 + F(0))\frac{1+\pi}{1+\pi^2+2F(0)\pi}\pi = \pi(1+\pi)/(1+\pi^2)$, and $\hat{\pi}\left(YYZZ; \bar{y}^{\mathsf{gl}*}_{\mathsf{S}} = 0\right) = 0$, and the RHS of (14) becomes $E\left[\widetilde{Z}\right]$ for $\bar{y}^{\mathsf{gl}*}_{\mathsf{S}} = 0$. 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 (13). That is, for any finite λ, $N^* \geq 2$.■

Proof of Proposition 6: Suppose $\lambda = 0$. Then, $W_{\mathsf{ia}}(0, \pi)$ is equal to $W(0, \pi)$ in case *one* agent reports to the centre that $y \geq \bar{y}^{FB}_{\mathsf{ia}}$ or $y < \bar{y}^{FB}_{\mathsf{ia}}$. 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) $\bar{\lambda}^{\mathsf{lo}}_{\mathsf{cl}} > \bar{\lambda}_{\mathsf{ia}}$ (see Propositions 1 and 4), and (ii) for all λ agents send influential information under cl *cum* gl. The truth of the final statement in the proposition has been verified numerically.■

Proof of Proposition 7: It follows from Propositions 2 and 4 that $\bar{\lambda}^{\mathsf{lo}}_{\mathsf{dl}} > \bar{\lambda}^{\mathsf{lo}}_{\mathsf{cl}}$ iff (17) holds. The existence of λ_6 then follows from the continuity of W in λ.■

Proof of Proposition 8: Consider cl, and suppose $N = 3$. We know $E\left[\widetilde{Z}|0 = a_0 \leq z \leq a_2\right] - a_1 = E\left[\widetilde{Z}|a_1 \leq z \leq a_3 = 1\right] - a_2$ from (11). If two becomes the maximum number of ranks, then $a_1 = 0$, and so this equality becomes $E\left[\widetilde{Z}|0 \leq z \leq a_2\right] = E\left[\widetilde{Z}\right] - a_2$. For any f_X, let $a_2^* < E\left[\widetilde{Z}\right]$ denote the unique value of a_2 satisfying this equality. Let $\mathbf{a}^*_{2/3} := (0, 0, a_2^*, 1)$. Hence, (11) and (12) become $\lambda\left[\hat{\pi}\left(YY; \bar{y}^{\mathsf{lo}}_{\mathsf{S}}, \mathbf{a}^*_{2/3}\right) - \hat{\pi}\left(YZ; \bar{y}^{\mathsf{lo}}_{\mathsf{S}}, \mathbf{a}^*_{2/3}\right)\right] = E\left[\widetilde{Z}\right] - a_2^*$ and $E\left[\widetilde{Z}|s^Y, s^Y, \bar{y}^{\mathsf{lo}}_{\mathsf{S}}\right] - \bar{y}^{\mathsf{lo}}_{\mathsf{S}} = E\left[\widetilde{Z}\right] - a_2^*$. There is a unique $\bar{y}^{\mathsf{lo}}_{\mathsf{S}}(\pi)$ that satisfies the latter equality. We can then use $\lambda\left[\hat{\pi}\left(YY; \bar{y}^{\mathsf{lo}}_{\mathsf{S}}(\pi), \mathbf{a}^*_{2/3}\right) - \hat{\pi}\left(YZ; \bar{y}^{\mathsf{lo}}_{\mathsf{S}}(\pi), \mathbf{a}^*_{2/3}\right)\right] = E\left[\widetilde{Z}\right] - a_2^*$ to find $\underline{\lambda}(\pi)$. For $\lambda \geq \underline{\lambda}(\pi)$, agents use at most two ranks. $\bar{\lambda}(\pi)$ is obtained from our numerical simulations. We checked the statement for $\pi \in [0.05, 0.95]$.■

Proof of Proposition 9: Fix λ, π, and f_X. Suppose $X_{1,1} = X_{2,1}$. A straightforward comparison of (8) and (14) 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 $y < \bar{c}^*_{\mathsf{D}} \leq z$, 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 $z < \bar{c}^*_{\mathsf{D}} \leq y$. However, for $(y, z) \in [0, \bar{c}^*_{\mathsf{D}}]^2$, both agents switch, while if $(y, z) \in [\bar{c}^*_{\mathsf{D}}, 1]^2$, both agents continue. In either case, the inferior technology is used at one site with probability one. Clearly, if $a_1^* = \bar{c}^*_{\mathsf{D}}$, then cl and dl would yield the same expected welfare. For given parameter values, they are, however, not the same. $\bar{c}^*_{\mathsf{D}}$ satisfies $\lambda\frac{\pi}{1+\pi} = E\left[\widetilde{Z}\right] - \bar{c}^*_{\mathsf{D}}$ (see (9)), whereas a_1^* satisfies $\lambda\frac{\pi}{1+\pi}4F(a_1^*)(1 - F(a_1^*)) = E\left[\widetilde{Z}\right] - a_1^*$ (see (15)). 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 $E\left[\widetilde{Z}\right] - \overline{c}_{\mathsf{D}}^*$, cl yields a higher expected welfare than dl.■

References

Austen-Smith, David. 1990. "Information Transmission in Debate." *American Journal of Political Science*, 34: 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*, ed. David Austen-Smith and John Duggan, Berlin: Springer-Verlag.

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

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 Equilibrium." 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." http://www0.gsb.columbia.edu/faculty/wdessein/.

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: 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.

Friebel, Guido and Michael Raith. Forthcoming (forthcoming). "Resource Allocation and Firm Scope." American Economic Journal: Microeconomics.

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: 287-335.

Hayek, Friedrich A., 1945. "The Use of Knowledge in Society." *American Economic Review*, 35: 519-530.

Gibbons, Robert. 1992. *A Primer in Game Theory.* Pearson Education, Harlow, UK.

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: 59-77.

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

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

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

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

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

Ottaviani, Marco and Peter Norman Sørensen. 2001. "Information aggregation in debate: who should speak first?." *Journal of Public Economics*, 81: 393-421.

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

Phelps, Charles E. 2000. "Information Diffusion and Best Practice Adoption." In *Handbook of Health Economics, Volume 1A*, ed. 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*, ed. 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: 878-906

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

Prendergast, Canice and Stole. 1996. "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning." *Journal of Political Economy*, 104: 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: 1257-1285.

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: 2817-2838.

Swank, Otto H. and Bauke Visser. Forthcoming. "Is Transparency to no Avail? Committee Decision-Making, Pre-meetings, and Credible Deals." Scandinavian Journal of Economics.

Thaler, Richard H. 2000. "Towards a Positive Theory of Consumer Choice." *Journal of Economic Behavior and Organization*, 1: 39-60.

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.