Filtered Social Learning*

Paul Niehaus†
UC San Diego

This Draft: February 9, 2011

Abstract

Knowledge-sharing is economically important but also typically incomplete: we "filter" our communication. This paper analyzes the consequences of filtering. In the model, homogenous agents share knowledge with their peers whenever the private benefits exceed communication costs. The welfare implications of this transmission mechanism hinge on whether units of knowledge complement, substitute for, or are independent of each other. Both substitutability and complementarity generate externalities; cheaper communication eliminates externalities in the former case, but not necessarily the latter. Complementary basic skills like numeracy catalyze technology adoption, and adoption may be path-dependent even when payoffs are certain and independent across agents.

JEL codes: D83, O31, O33

Keywords: Communication, Social Learning, Technology Adoption

*I thank Ed Glaeser and Sendhil Mullainathan for encouragement and guidance throughout this project. I benefited from conversations with Nageeb Ali, Ryan Bubb, Eric Budish, Daniel Carvalho, Vince Crawford, Michael Faye, Drew Fudenberg, Jerry Green, Justin Ho, Michael Kremer, Markus Möbius, Joel Sobel, and Sandip Sukhtankar, and from the input of seminar participants at Harvard, Cornell, LSE, Warwick, University of California at San Diego, University of Wisconsin at Madison, and Dartmouth College. I also wish to thank the editor, Philip Reny, and two anonymous referees for suggestions that substantially improved the paper, and Aislin Bohren for outstanding research assistance. I acknowledge financial support from a National Science Foundation Graduate Research Fellowship. All errors are my own.

†Department of Economics, University of California at San Diego, 9500 Gilman Drive #0508. pniehaus@ucsd.edu.

1 Introduction

Knowledge spillovers play an important role in economic activity. Word-of-mouth helps shape decisions about consumption (Moretti, 2011) and saving (Duflo and Saez, 2003). Peer-to-peer knowledge transmission drives the adoption of new technologies (Foster and Rosenzweig, 1995; Kremer and Miguel, 2007; Conley and Udry, 2010) and the creation of new ones (Jaffe et al., 1993). Besides explaining micro-economic behavior, these phenomena are predicted to have powerful aggregate implications. For example, knowledge spillovers are seen as determinants of economic growth (Romer, 1986; Lucas, 1988) and of spatial equilibrium (Black and Henderson, 1999; Glaeser, 1999). Understanding the (in)efficiency of spillover processes thus appears to be an important task.

Micro-theorists have modeled spillovers in two ways. The observational learning approach restricts agents to observing each others actions and/or payoffs; one has of course the option of interpreting "observation" as "communication".[1] An alternative approach models communication explicitly but with exogenous restrictions on what people do or do not communicate, rather than a full specification of costs and benefits.[2] This has two important consequences. First, it limits welfare analysis – though most would agree that externalities are likely to arise and to be economically important. Second, one suspects that under reasonable cost specifications the behavior posited would be inefficient even locally, i.e. from the point of view of pairs of communicating agents. For example, a receiver in the persuasion bias framework (DeGroot (1974), Demarzo et al. (2003)) would want to learn whether he is receiving independent or correlated signals from his various sources, but does not. In such cases it is unclear to what extent adverse outcomes, such as failure to adopt a beneficial technology, are due to local contractual failings or bounded rationality as opposed to communication externalities.

This paper proposes a different approach in which the costs of communication are modeled explicitly. Agents are endowed with discrete units of knowledge and must weigh the costs and benefits of communicating each unit. I call this process *filtering*. Some filtering is socially valuable because communication takes time and effort; simply transferring all of our knowledge to each other would be wasteful. But filtering may also impose social costs if particular units of knowledge do not reach the people who need them. When and why will this be the case?

[1]See Banerjee (1992), Bikhchandani et al. (1992), Banerjee (1993), Ellison and Fudenberg (1993), Besley and Case (1994), Bala and Goyal (1998), and Acemoglu et al. (2008).

[2]See DeGroot (1974), Jovanovic and Rob (1989), Ellison and Fudenberg (1995), Glaeser (1999), Demarzo et al. (2003), Banerjee and Fudenberg (2004), Kondor and Ujhelyi (2005), Calvó-Armengol and de Martí (2007), Jackson and Golub (2010), and Kremer and Miguel (2007). Acemoglu et al. (2010) are exceptional in that they study endogenous choices about whether and to whom to communicate. They focus on conditions for asymptotic learning rather than on welfare comparisons, however, and on learning about a single parameter rather than on complementary or substitutable knowledge (as this paper will).

I study these questions under the assumption that communication is *locally efficient*: agents share knowledge if and only if the benefits to the learner exceed the costs of communication. This serves two purposes. First, it helps clarify what kinds of inefficiency are due to externalities, as opposed to local contractual failings. Second, I argue that it approximates several real-world settings of interest. A great deal of social learning takes place between friends and acquaintances motivated to help each other. Knowledge sharing may also occur in the workplace where co-workers engaged in team production share knowledge in order to raise group productivity. The settings I have in mind are explicitly non-strategic, and the model thus complements recent work on sequences of strategic communication (Stein, 2008; Ambrus et al., 2010).

The requirement that communication choices respond to the local costs and benefits of knowledge-sharing imposes some discipline on aggregate behavior; the question is, how much? I study this question in a sequential setting where agents speak with predecessors who have been in similar situations. All agents share a common value function for knowledge. The utility of this approach is that it permits analysis of fully general payoff functions describing the returns to knowledge. For example, one can think coherently about disparate units of knowledge such as numeracy and the ability to identify insect eggs within the same framework. (Exactly why the latter is useful will be made clear.)

Section 3 characterizes the efficiency of learning over finite sequences of conversations. I show that efficiency hinges on whether different kinds of knowledge complement, substitute, or are independent of each other. In the independent case communication is efficient from a social point of view, even though each communication decision takes into account only local costs and benefits. The reason is that under independence an agent's marginal valuation of knowledge does not depend on what he already knows, making the marginal net returns to communication the same everywhere.

This symmetry property breaks down when different kinds of knowledge either substitute for or complement each other, because in these cases each agent's marginal valuation for knowledge varies depending on what he already knows. For example, knowledge about different vendors selling similar goods is substitutable, and so communication externalities can hinder price competition. The units of knowledge required to implement a new technology are complementary, so adoption may be inefficiently slow even when the returns to the new technology are known.[3] These externalities always involve "under-communication" relative to what a planner would have chosen, though interestingly there can be situations where lack of communication at one juncture is triggered by "over-communication" at an earlier one.

Since the social benefits of communication sometimes exceed the private benefits there

[3]Rogers (2003) opens his review thus: "Many innovations require a lengthy period of many years from the time when they become available to the time when they are widely adopted" (1).

may be scope for welfare-improving subsidies. Directly subsidizing communication would be difficult, but governments often fund interventions that lower communication costs: for example, construction of communication and transportation infrastructure and training in skills like reading, writing, and language. Whether there are external returns to these investments is an active area of research.[4] I therefore examine whether the model can rationalize them by characterizing the welfare effects of cheaper communication.

The answer is "maybe." When knowledge is substitutable welfare losses are bounded above by the costs of communicating, so that lowering those costs sufficiently must eliminate externalities. When knowledge is complementary, however, this need not hold, because agents may rationally choose not to incur even an arbitrarily low communication cost if a complementary unit of knowledge is not available. The argument for subsidy is thus less immediate for learning processes like technology adoption and innovation where complementarities are generic.[5]

Section 4 examines whether communication externalities are eventually overcome or whether they may also hinder asymptotic learning, along with implications for technology adoption. Previous work has studied asymptotic learning about an unknown parameter and has emphasized factors such as agents' rules-of-thumb (Ellison and Fudenberg, 1993), the scope and structure of social interactions (Ellison and Fudenberg, 1995; Bala and Goyal, 1998; Acemoglu et al., 2008, 2010), and the strength of signals (Smith and Sorensen, 2000; Banerjee and Fudenberg, 2004). In the filtering model the main impediment to asymptotic learning is complementarity: if knowledge is too disaggregated initially then it may never begin to accumulate.

As this suggests, knowledge can be catalytic in the sense that it raises the value of other skills and thus accelerates social learning. Development economists since at least Nelson and Phelps (1966) have argued that human capital matters as "a factor that facilitates diffusion" as well as "a factor of production" (Benhabib and Spiegel, 2005). With filtered communication these two functions coincide: human capital facilitates diffusion precisely because it complements new techniques in the production function. Section 4.3 illustrates this with an example of technology diffusion in which a basic skill, numeracy, complements a new farming technique. Because of filtering, (1) adoption is asymptotically incomplete and (2) there are socially increasing returns to the basic skill. This may help explain evidence from diffusion research that better-educated individuals, as well as individuals with more access to new knowledge, most influence their peers (Rogers and Shoemaker, 1971).

Finally, I examine whether filtering can lead to multiple stable steady-states, and in

[4] See Gramlich (1994) and more recently Straub (2008) for reviews on returns to infrastructure, Heckman and Carneiro (2003) and Moretti (2004) on external returns to education.

[5] Subsidizing skills needed to adopt or innovate may still have high returns in those cases, however. Section 4.3 examines this.

particular to lock-in to either of two competing technologies. Lock-in is often observed in practice; for example, Rogers and Kincaid (1981) document how women in each of 25 Korean villages tended to converge to use of the same method of birth control – the pill in some villages, IUDs in others, vasectomy in yet others. Several papers have shown that lock-in may arise when there is uncertainty about payoffs; indeed, even an isolated agent experimenting with uncertain alternatives may rationally settle on an inferior one. It is also well-known that if the payoffs to adopting a technology depend on the number of other users (as with telephones or email) then society may converge to a Pareto-inferior equilibrium. Section 4.4 shows that filtering can generate lock-in even with independent and certain payoffs. The necessary condition is that some knowledge be substitutable, which is typically the case in technology choice problems as skills specific to one technology substitute for skills specific to the other. The mechanism is intuitive: as knowledge about one technology accumulates this depresses incentives to communicate about the alternative, eventually generating lock-in even if the alternative is intrinsically preferable.

Some discussion of avenues for further research is in Section 5, the conclusion.

2 A Model of Filtered Social Learning

2.1 Primitives

For a general analysis of issues like complementarity and substitutability we need a flexible way of representing knowledge and its value. Let there be n discrete units of knowledge. Each unit can be thought of as a technique for accomplishing something useful; for example, one unit of knowledge could be knowing how to work out fractions. Agents either know or do not know each unit of knowledge, and so we can represent their knowledge as an element of $K \equiv \{0,1\}^n$. k_t is agent t's *knowledge endowment*; $k_t^i = 1$ if agent t knows the ith unit of knowledge, and $k_t^i = 0$ otherwise. The utility of this setup is that it can accommodate arbitrary kinds of incommensurable knowledge – for example, one unit of knowledge might be knowing how to identify insect eggs, while another could be knowing how to prove Kakutani's theorem. The framework lets us to think coherently about both together.

I use standard notation defined on K: the join (or coordinate-wise maximum) operator $\vee$ and the meet (coordinate-wise minimum) operator $\wedge$; the partial product order $\precsim$ which ranks $k \precsim \hat{k}$ if $k^i \leq \hat{k}^i$ for all i and its implied strict order $\prec$; the Euclidean norm $|\cdot|$; and a difference operator $\setminus$ defined by $k \setminus k' = k - (k' \wedge k)$. e_b is the unit vector with a 1 in the bth spot and zeros otherwise, and $\overline{k} = (1, \ldots, 1)$ is an n-vector of ones. I say that endowment k contains unit of knowledge $b \in \{1, \ldots, n\}$ if $k^b = 1$.

Assumption 1 (Homogeneity)**.** *All agents share a common value function* $V : K \to \mathbb{R}$ *over knowledge endowments.*

Agents are homogenous in that they share common preferences over knowledge endowments. This rules out situations where different farmers need to learn different things because their asset holdings differ (see e.g. Munshi (2004)). The value of this approach is that it focuses attention on generic issues in knowledge aggregation. Of course, type heterogeneity will play an important role in many applied settings. Understanding the efficiency of communication among homogenous agents can help predict incentives for assortative matching.

Assumption 2 (Monotonicity)**.** *If* $k' \succsim k$ *then* $V(k') \geq V(k)$.

The value function V captures the payoff that agents get from maximizing an underlying profit or utility function. Knowing more expands the feasible set of actions, so that a larger endowment of knowledge is always weakly beneficial. Its value may at times be less than the cost of communicating it, however, leading to filtering.

2.2 Interaction and Conversation

I study communication in a sequential setting. A countably infinite set of agents indexed $t = 1, 2, \ldots$ each deal with the same situation in turn. Agent t's initial endowment of knowledge $\tilde{k}_t$ is drawn independently according to the c.d.f. F on K, reflecting things she has learned from formal education, exposure to advertising, observation, learning-by-doing, and so forth. Leaving this endowment exogenous sharpens our focus on issues to do with the transmission, rather than the initial acquisition, of knowledge. I remark on some implications of endogenous acquisition below.

The timing of play is simple: at time t agent t speaks with one of her antecedents $a(t)$, where $a(t) < t$, and then takes an action and receives a once-and-for-all payoff V. Let the probability $q(t, \tau)$ that $a(t) = \tau$ satisfy $0 \leq q(t, \tau) \leq 1$ for all (t, τ), $q(t, \tau) = 0$ for $\tau \geq t$, and $\sum_\tau q(t, \tau) = 1$ for all t.[6] Also, let $A(t) = \{a(t), a(a(t)), \ldots, 1\}$ be the set of all of t's antecedents – this is the set of agents whose knowledge it would be feasible for t to acquire.

This structure embeds a variety of commonly studied interaction processes, such as workers learning from their predecessors on the job, children learning from their parents, farmers learning from more experienced neighbors, etc. The exogenous pattern of social interactions is intended to capture the idea that opportunities to communicate are often driven by cost considerations or by factors at least partly orthogonal to the demand for knowledge.[7] While endogenizing the search for communication partners is also an

[6] Formally, $q(t, \tau)$ defines a random graph model (Bollobás, 2001).

[7] For example, Marmaros and Sacerdote (2006) show how friendships form due to coincidental proximity.

important task for the theory of social learning, this paper limits itself to endogenizing communication within a fixed interaction structure.

2.3 Filtering through Conversation

Both formal empirical work (e.g. Conley and Udry (2010)) and casual observation make it clear that knowledge spillovers are typically incomplete. Sinha and Mehta (1972) provide an elegant illustration: they measured the incompleteness of communication between a set of "contact farmers," who were given information about irrigation techniques, and their peers. When they surveyed the secondary recipients – farmers who had acquired some of the new knowledge in conversations with the contact farmers – they found that on average only 28% of the knowledge had spilled over. They also found that the secondary recipients who learned the most were those who spoke to the primary recipients who learned the most. These facts illustrate how communication networks "filter" knowledge-sharing and thus determine what knowledge each individual has access to.

The main innovation of this paper is to endogenize the filtering process as the solution to a well-defined optimization problem. Suppose an agent r (the "receiver") speaks with an agent s (the "sender"); in our model the receiver at time t will be agent t and the sender will be $a(t)$. The sender's knowledge endowment is k_s while the receiver's is $\tilde{k}_r$ at the beginning of the conversation and k_r at the end of it. I require that the receiver does not loose any knowledge during the conversation ($k_r \succsim \tilde{k}_r$) nor learn anything that agent s did not know ($k_r \precsim \tilde{k}_r \vee k_s$). Conversation thus serves purely to communicate knowledge, not to create it (in contrast with Jovanovic and Rob (1989) and Stein (2008)). Define $K_r \equiv \{k : \tilde{k}_r \precsim k \precsim \tilde{k}_r \vee k_s\}$ the set of feasible final knowledge endowments for r; the economic problem is to select one from among them. Any theory of communication must take a stance on this issue, if only implicitly.

Here I suppose that the agents choose to communicate what is *locally optimal*, trading off the private benefits of communication against its costs. The private benefit of communication is the increase in the instantaneous payoff the receiver will obtain. Costs are represented as a fixed vector $c = (c^1, \ldots, c^n)$ with c^b representing the cost of transmitting the bth unit of knowledge. I then define a map $C : K \times K \to K$ from the parties' initial knowledge endowments to the receiver's final knowledge endowment:

Assumption 3 (Filtering)**.**

$$C(k_s, \tilde{k}_r) \in \arg\max_{k \in K_r} V(k) - c \cdot (k - \tilde{k}_r) \tag{1}$$

As a tie-breaking rule if there are multiple maximizers then one of those with the smallest norm is chosen.

In other words, communication maximizes benefits net of costs. The costs c in this definition reflects real costs of communicating: primarily the time and effort involved. Both parties bear these costs to some extent, but since communication is assumed efficient we need not take a stance on their exact division.[8]

Assumption 3 implies that agents align their incentives locally but do not take into account the broader social consequences of their choices. I adopt this approach for two reasons. First, as noted in the Introduction, existing work on communication assumes communication that is inefficient *locally*, i.e. from the point of view of the communicating parties. Introducing local efficiency thus plays a useful logical function: it helps us understand what kinds of inefficiencies are intrinsic to the process of aggregation, as opposed to those that simply mirror local inefficiencies.

The second function of Assumption 3 is to approximate a variety of real-world situations of interest. It seems plausible that a good deal of communication is motivated by altruism: people share knowledge in order to help their friends and acquaintances. Under this interpretation the sender directly internalizes the receiver's utility gains from learning when deciding what to communicate. Alternatively, local incentives may arise in team production settings. If the sender and receiver work in a team and the sender's compensation varies one-for-one with the receiver's output then again the sender will internalize the immediate costs and benefits of sharing knowledge. In both these cases forward-looking receivers might like the sender to share additional knowledge which they could then share with the agents they subsequently speak with, but are unable to contract with the sender to do this. These interpretations rest on the idea that one cannot make "a thoroughly appropriable commodity of something so intangible as information" (Arrow, 1962). Such contracts would be prohibitively costly to write or enforce, and the very act of describing the desired transaction may give away the knowledge itself – Arrow's "paradox of information."[9]

Interestingly, however, one can also interpret Assumption 3 as an equilibrium outcome when local knowledge-sharing is *fully* contractible; the wedge in this case is holdup. Suppose the receiver decides what knowledge to purchase from the sender, but cannot specify re-sale terms in advance with the agents she will subsequently speak with. Then those agents will hold up part of her sunk investment in acquiring knowledge from the sender. The problem becomes more acute as the relative bargaining power of future buyers increases, and in the limit one recovers Assumption 3.[10]

A final comment on Assumption 3 concerns the actual process of optimization. In

[8]Dewatripont and Tirole (2005) also study costly communication but in a single exchange and between agents with imperfectly aligned preferences.

[9]In this vein, Jovanovic and Nyarko (1995) study the inter-general transmission of knowledge through a market subject to adverse selection.

[10]Formally, consider a sender s and receiver r deciding on a final knowledge endowment k_r, and suppose

general, finding a locally optimal communication choice will require that someone know what units of knowledge are contained in both agents' endowments. If at the beginning of the conversation neither agent knows what the other knows then we are implicitly assuming that they find out during the course of the conversation. For example, the receiver might say "I know how to do this but I'm having trouble with that," to which the sender responds by explaining "that." This seems reasonable in settings where it is much easier to describe what one does or does not know than to actually transfer that knowledge. For example, asking how to prove Kakutani's theorem takes much less time than explaining the proof. Assumption 3 best applies to situations such as these. Introducing additional costs of communicating about what to communicate about leads to models of the conversation process itself, which are interesting but beyond the scope of this paper.

3 Communication Externalities

This section characterizes the relationship between efficiency and the value function (V) – in particular, whether different kinds of knowledge complement, substitute for, or are independent of each other. I define welfare over groups of agents additively:

Definition 1 (Welfare)**.** *Fix a sequence of agents* $1, \ldots, t$*, a sequence of endowments* $\tilde{k}^t = \tilde{k}_1, \ldots, \tilde{k}_t$ *on support*(F)*, and a mapping* $a(2), a(3), \ldots, a(t)$ *from agents to the predecessor each speaks with. A sequence* $k^t = k_1, \ldots, k_t$ *on* K *is feasible iff* $\tilde{k}_\tau \precsim k_\tau \precsim \tilde{k}_\tau \vee k_{a(\tau)}$ *for all* τ*. For any such feasible sequence social welfare is*

$$W(k^t, \tilde{k}^t) = \sum_{\tau=1}^{t} V(k_\tau) - c \cdot (k_\tau - \tilde{k}_\tau) \tag{2}$$

Call the sequence generated by $k^t = C(k_{a(t)}, \tilde{k}_t)$ the equilibrium outcome, and say that learning is efficient over the sequence $1, \ldots, t$ if the equilibrium outcome maximizes

that the sender bears a share α of the communication costs. Their payoffs, gross of transfers, are

$$U_s = -\alpha c \cdot (k_r - \tilde{k}_r)$$
$$U_r = V(k_r) - V(\tilde{k}_r) - (1-\alpha)c \cdot (k_r - \tilde{k}_r) + R(k_r) - R(\tilde{k}_r)$$

where $R(k)$ is the receiver's expected payoff from re-selling units of knowledge from the knowledge endowment k. If the receiver has full bargaining power then for any k_r he will propose a monetary transfer $T(k_r) = \alpha c \cdot (k_r - \tilde{k}_r)$, just compensating the sender for her time and effort, so that her net payoff is 0 irrespective of k_s. As r's future interactions have a symmetric structure it must also hold that $R(k) = 0$ for all k. r is therefore left to solve

$$\max_{k_r} V(k_r) - V(\tilde{k}_r) - (1-\alpha)c \cdot (k_r - \tilde{k}_r) - \alpha c \cdot (k_r - \tilde{k}_r)$$

which differs from Equation 1 only by the constant term $V(\tilde{k}_r)$.

Figure 1: Costly Communication

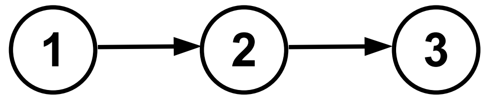

(2) among all feasible length-t sequences of elements of K. If not, say that an externality occurs.[11]

Since one might reasonably expect the model to generate externalities, I start with a case where it does not. Consider a sequence of three agents, as depicted in Figure 1. Suppose that these are farmers and that each in turn needs to deal with a pest problem. (I will use pest control as a running illustration to illustrate how the same applied problem may have different efficiency properties depending on the kind of knowledge in play.) Suppose Agent 1 knows something that the others do not about pest control. Further suppose that the value of this knowledge is *independent* of what else they know. Call V the (unconditional) incremental value of the knowledge, and c the cost of communicating it. The socially optimal communication scheme depends on whether $2V > 2c$. If so then it is socially optimal for agent 1 to share his knowledge with 2 and 2 with 3; otherwise "no communication" is socially optimal. On the other hand, it is *locally* optimal for agent 1 to share what he knows with agent 2 if and only if $V > c$, and if this holds then it will also be locally optimal for agent 2 to relay it to 3. The conditions for private and social optimality thus coincide exactly.

This reasoning extends directly to longer sequences of conversations and to any number of units of knowledge whose values are all independent of each other:

Proposition 1. *If V is linear (can be written $V(k) = \sum_{b=1}^{n} V_b k^b$) then, in equilibrium, learning over any finite sequence of agents is socially optimal.*

Proof. By assumption we have for any b either $V(k \vee e_b) - V(k) > c^b$ for all k not containing b, or $V(k \vee e_b) - V(k) \leq c^b$ for all k. Call the set of bits satisfying the former condition B_1 and the latter B_0. Bits $b \in B_1$ will always be communicated from senders who have them to receivers who do not, so in equilibrium $k_t^b = \max_{\tau \in A(t)} k_\tau^b$. Bits $b \in B_0$ will never be communicated and so $k_t^b = \tilde{k}_t^b$. I next show that this is the socially optimal outcome. Consider any feasible scheme in which for some $b \in B_1$ and some t, $k_t^b < \max_{\tau \in A(t)} k_\tau^b$. Then these exists some $\tau \in A(t)$ for which $k_{a(\tau)}$ contains b but k_τ does not. Then the (feasible) scheme in which k_τ contains b is welfare-improving, so the

[11]This definition implicitly sets a high bar, in that eliminating externalities might require knowing what knowledge all agents are endowed with. However, under Assumption 3 all agents would behave identically even if they knew about everyone else's endowments. Welfare losses in this section are therefore best thought of as resulting from an incentive problem.

original scheme cannot be socially optimal. Similarly, if for some $b \in B_0$ and some t we have $k_t^b > \tilde{k}_t^b$ then we can construct a feasible welfare-improving alternative in which k_t does not contain b. Hence the equilibrium outcome is a sequence of endowments that matches the socially optimal one, and clearly this is achieved at minimum cost. □

This efficiency result is driven by symmetry across agents. Constant returns implies that the marginal value of each unit of knowledge is not affected by what an agent initially knows, and because all agents share the same value function this marginal value is identical across agents. Along with the assumption that costs are the same across agents this implies that each agent shares knowledge if and only if it is efficient for all other agents to do so.

Previous work on social communication has not emphasized communication externalities, so comparisons between this result and earlier models inevitably depend somewhat on how one re-interprets their assumptions. It does seem fair to say, however, that efficiency is not a common property. One area in which the relationship between communication and efficiency is better-understood is work on strategic experimentation, which has argued that agents tend to acquire too little costly information because they do not internalize its value to society (Foster and Rosenzweig, 1995; Bolton and Harris, 1999; Bramoulle and Kranton, 2007). Interpreted through the lens of Proposition 1 we see that this literature's key assumption is an asymmetry between the cost of initially acquiring knowledge and the costs of subsequently diffusing it. For example, a typical assumption might be that knowledge, once discovered, becomes freely available to everyone within a particular village. Proposition 1 also provides one justification for this assumption: the combination of (i) independently useful knowledge and (ii) local incentives to share knowledge is sufficient for widespread diffusion.

3.1 Substitutable Knowledge

Proposition 1 highlights the role of symmetry in the model, in which only knowledge endowments vary. Of course, heterogeneity in knowledge endowments is necessary to make a social learning model interesting. This brings us to an interesting question: can the very feature that makes social learning possible also cause it to break down?

One way in which this can happen is if different kinds of knowledge substitute for each other. For example, suppose a series of farmers face pests on their crops at different times. There are various methods of dealing with pests and the farmers initially know different ones, where knowing a method means knowing what inputs to purchase, where to purchase them, how to apply them, when to apply them, etc. For simplicity suppose that all methods are equally effective but that they require different expenditures on labor and other inputs, so that farmers will be interested in learning the cheapest methods.

Figure 2: Dealing with Pests

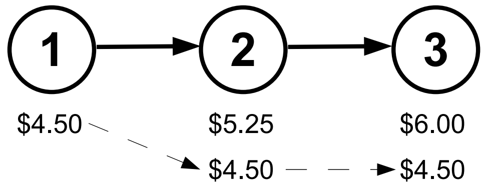

Explaining any method to a peer costs $1.

Figure 2 depicts a sequence of three agents with the cost of the cheapest pest-control method each agent is initially endowed with recorded beneath him. Communication in this example evolves as follows. First, Agent 2 speaks with Agent 1. Agent 1 knows a cheaper way of dealing with pests than Agent 2, but because sharing this knowledge would cost $1.00 and would only save Agent 2 $0.75 it is not locally optimal to share it. Next Agent 3 speaks with Agent 2, and similarly it is not locally optimal for Agent 2 to teach him anything. It would have been feasible, however, for Agent 1 to teach the $4.50 method to Agent 2 and for Agent 2 to then teach it to Agent 3, which would have saved him $(\$6.00 - \$4.50) - \$1 = \$0.50$. Since this gain more than offsets Agent 2's $0.25 net loss from learning the $4.50 method, an externality has occurred: total costs in this economy are higher than they could have been. The root problem here is the imperfect substitutability of methods: Agent 2's knowledge of a $5.25 method for controlling pests lowers her marginal valuation of Agent 1's knowledge below agent Agent 3's marginal valuation.

An interesting feature of this example is that Agent 3 – and society as a whole – would be better-off if Agent 2 knew *less*. Suppose that in the scenario just described Agent 2 knew both the $5.25 method and the $6.00 method, and then consider a revised scenario in which Agent 2 knows only the $6.00 method; in this case she and Agent 3 would learn the $4.50 method, and the total cost incurred by the two of them will fall. This suggests that filtering may have interesting implications for competition: a firm that enters the market with a product that costs $5.25 will disrupt social learning about the $4.50 product, exacerbating "business-stealing" motives for entry (Mankiw and Whinston, 1986).

Another way to look at this example is that Agent 3 suffers from speaking with a peer with a disparate knowledge endowment. Typically one expects agents to benefit most from speaking with peers whose knowledge endowments differ the most from theirs, since those differences create a kind of gain from trade (Glaeser and Sunstein, 2007). But in the presence of filtering peers who know different things will also learn different things and in particular may not always learn the things that one would want them to, damping the benefits of heterogeneity.

Note also that the inefficiency in this example is somewhat fragile. If the \$4.50 method were slightly cheaper, so that it improved on the \$5.25 method by more than \$1, then the externality would vanish. Similarly if the cost of communicating were slightly lower then Agent 2 would learn the \$4.50 method, again eliminating the externality. This suggests a role for policy: anticipating problems of the sort just described, a planner might subsidize communication infrastructure (e.g. mail service) or the acquisition of communication skills (e.g. reading, writing). The next proposition generalizes this idea, showing that lowering communication costs far enough will eliminate externalities for *any* submodular value function.

Proposition 2. *If V is submodular then in equilibrium, and for every realization of agents' endowments, the average welfare loss among any finite set of agents is at most* $\sum c^b$.

Proof. I will show that no single agent's welfare is lower in equilibrium than in the optimum scheme by more than $\sum c^b$, which implies that this must also hold on average. Fix any agent t and consider the sequence of his antecedents $a(t), a(a(t)), \ldots, 1$; for calculating t's welfare there is no loss of generality in re-labeling agents so that $t = |A(t)| + 1$ and $a(\tau) = \tau - 1$ for $\tau \leq t$. Consider then the sequence of initial endowments $\tilde{k}_1, \ldots, \tilde{k}_t$ generating final endowments $k_1, \ldots, k_t$, and let $\hat{k}_1, \ldots, \hat{k}_t$ be a feasible alternative sequence of final endowments. The difference G in agent t's welfare is

$$\begin{aligned} G &\equiv V(\hat{k}_t) - V(k_t) - c \cdot ((\hat{k}_t - \tilde{k}_t) - (k_t - \tilde{k}_t)) \\ &\leq V(\overline{k}_t) - V(k_t) + c \cdot (k_t - \tilde{k}_t) \end{aligned}$$

where $\overline{k}_t = k_1 \vee \ldots \vee k_t$. Expand

$$V(\overline{k}_t) - V(k_t) = \sum_{\tau=2}^{t} V(k_{\tau-1} \vee \ldots \vee k_t) - V(k_\tau \vee \ldots \vee k_t)$$

Each of these marginal differences represents the marginal value of knowledge that agent $\tau - 1$ has but none of agents $\tau, \ldots, t$ have. Let $d_\tau = k_{\tau-1} \setminus (k_\tau \vee \ldots \vee k_t)$ be an endowment containing that knowledge. Then

$$\begin{aligned} V(k_{\tau-1} \vee \ldots \vee k_t) - V(k_\tau \vee \ldots \vee k_t) &= V(d_\tau \vee k_\tau \vee \ldots \vee k_t) - V(k_\tau \vee \ldots \vee k_t) \\ &\leq V(d_\tau \vee k_\tau) - V(k_\tau) \\ &\leq c \cdot d_\tau \end{aligned}$$

where the first inequality follows from submodularity and the second from the local optimality of conversations (Assumption 3). Summing up yields

$$\begin{aligned} G &\leq c \cdot (d_2 \vee \ldots \vee d_t) + c \cdot (k_t - \tilde{k}_t) \\ &\leq \sum c^b \end{aligned}$$

as desired. □

The intuition here is that if the nth unit of knowledge does not reach an agent t who demands it then it must be that some intermediate agent t' knew some other thing which was a close substitute, where "close" is defined relative to communication costs. If agent t does not learn that other thing either then the same argument applies, and so on. One can construct examples (available on request) in which average welfare losses approximate $\sum c^b$ arbitrarily closely, so this upper bound is tight. Note that Proposition 2 holds for every realization of endowments and so the same bound must also hold in expectation ex-ante, before endowments are drawn.[12]

Proposition 2 implies that if knowledge is substitutable then public investments in lowering communication costs can be welfare-enhancing. For example, learning about alternatives in a competitive market for a good like pest control need not be efficient, but interventions that lower the cost of sharing information about alternative pest control products can reduce and eventually eliminate the inefficiencies. These interventions could be specific to the market or could be general-purpose, as training in communication skills or provision of communication infrastructure.[13]

3.2 Complementary Knowledge

Knowledge about different ways of accomplishing the same task is substitutable. But sometimes accomplishing a task requires knowing several things. For example, writing to a colleague requires knowledge both of how to write and of his address. Treating an illness requires both skill in diagnosis and knowledge of the available treatments. Conducting empirical research requires both conceptual knowledge of empirical techniques and practical knowledge about how to implement them using statistical software. These are all examples of situations in which knowledge is complementary: knowing one thing raises the marginal value of learning another.

[12] While linear costs are a natural base case, one can also imagine situations in which costs are supermodular (due to increasing marginal costs of time) or submodular (due to inter-relationships between the knowledge being communicated). Proposition 2 can be generalized to these cases, replacing c^b with a bound on the marginal cost of communicating the bth unit of knowledge (available on request). It is the submodularity of V that is crucial.

[13] Jensen (2007) shows how lower communication costs improved market efficiency in India fisheries, though how much of these gains were due to reduced externalities cannot be ascertained.

Figure 3: Integrated Pest Management

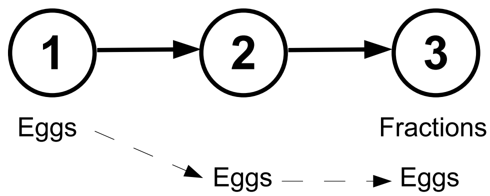

To illustrate the implications of complementary knowledge I take as an example the adoption of Integrated Pest Management (IPM), a skill-intensive set of practices designed to reduce pest loads while also reducing farmers' use of potentially toxic fertilizers. Winarto (2004) provides a vivid description of IPM adoption among Javanese farmers. As an abstracted version of her account, suppose that to perform IPM farmers must first count the insect eggs growing on their crops, distinguishing the eggs of harmful pests (e.g. rice stem-borer) from those of helpful natural predators (e.g. spiders), and then perform some calculations to determine whether or not spraying pesticide will be profitable. Spraying pesticide is beneficial when there are many pests and few predators, and wasteful otherwise.

Figure 3 depicts an example of a sequence of conversations about IPM where the complementary units of knowledge are initially dispersed. Agent 1 knows how to identify insect eggs but is not numerate, nor is Agent 2, and so it is not worthwhile for him to teach her what he knows. Given this, Agent 2 has nothing to teach Agent 3. It would have been feasible, however, for Agent 2 to relay Agent 1's knowledge about insect eggs to Agent 3, in which case Agent 3 could have adopted IPM. Adoption is thus delayed and welfare potentially decreased because of communication externalities.

The externality in this example has a key feature that distinguishes it from the example with substitutes: the equilibrium pattern of communication does not depend sensitively on the costs of communication or on the profitability of IPM. All that matters is that neither of the two units of knowledge required to perform IPM is of much use on its own. This suggests that lowering communication costs may not be as effective at controlling externalities here as it was in the submodular case. Proposition 3 proves this by example:

Proposition 3. *Fix $n \geq 2$ and the costs c of communication. For any $L > 0$ there exists an initial distribution F, interaction process $q(t, \tau)$, and supermodular V such that over some fixed sequence of agents the expected welfare loss exceeds L.*

Proof. Let $q(t, \tau)$ be 1 for $\tau = t - 1$ and 0 otherwise. Let $V(k) = \gamma k^1 k^2 + \hat{V}(k^3, \ldots, k^n)$ with $\hat{V}$ an arbitrary supermodular function and $\gamma > 2c^1$. Let F be such that with positive probability the first three agents are endowed with a sequence $\tilde{k}_1, \tilde{k}_2, \tilde{k}_3$ such that $\tilde{k}_1 = (1, 0, \tilde{k}_1^3, \ldots, \tilde{k}_1^n)$, $\tilde{k}_2 = (0, 0, \tilde{k}_2^3, \ldots, \tilde{k}_2^n)$, and $\tilde{k}_3 = (0, 1, \tilde{k}_3^3, \ldots, \tilde{k}_3^n)$. Let ϕ be the probability that such a sequence arises under F. Following any such sequence it must be that $k_3 = (0, 1, k_3^3, \ldots, k_3^n)$ regardless of γ even though $(1, 1, k_3^3, \ldots, k_3^n)$ was feasible at an additional communication cost of $2c^1$. The expected welfare loss over this sequence is thus at least $\phi \cdot (\gamma - 2c^1)$, which exceeds L if $\gamma > L/\phi + 2c^1$. □

The key point here is that one can continuously increase a parameter of the value function while holding communication costs fixed (and thus not merely scaling up the entire problem). This would be trivially possible in models where communication choices are exogenous and thus unresponsive to the value of knowledge. Here it is possible even though the local optimality condition ensures that more valuable knowledge is, in general, more likely to be communicated.

For policy this implies that interventions that indirectly subsidize communication by lowering communication costs (such as investments in communication infrastructure or communication skills) may not mitigate communication externalities when knowledge is complementary. Intuitively the problem is that agents will not bother incurring even a small communication cost to share knowledge if a key complementary skill or idea is missing. In these cases more direct subsidies to communication itself, including coordinating more centralized methods of communication, may be necessary. Subsidizing the acquisition of complementary knowledge might also be effective; Section 4.3 examines this idea in slightly more depth.

While adoption in this example hinged on the acquisition of complementary skills, similar principles would apply to communication about the *returns* to a new technology. Suppose the technology has two distinct advantages over its predecessor – for example, IPM is capital-saving and reduces environmental health threats. If knowing both of these advantages would convince a farmer to adopt, while knowing either one on its own would be insufficient, then they are complements. Locally optimal communication in such an environment would again lead to inefficient delays in adoption.

Complementary knowledge features in processes of innovation as well as adoption. Farmers themselves often innovate by combining traditional knowledge with new ideas obtained from extension services (Bentley, 2006), and more generally new ideas typically build on or combine existing ones (Jacobs, 1969; Feldman, 1994; Holton, 1998). Mokyr (1990) provides some salient examples in his discussion of punctuated technological change:

> “Lewis and Paul’s 1740 invention of the roller replacing human fingers as a yarn-twisting device had to wait until it was complemented a quarter-century later by Arkwright’s relatively marginal but crucial insight to use two (instead of one) sets of rollers. Bessemer’s invention of the converter would have been

useless had it not been for Robert Mushet's addition of spiegeleisen (an alloy of carbon, manganese, and iron) into the molten iron as a recarburizer."

Weitzman (1998) provides additional examples to motivate a model of productivity growth in which existing ideas serve as complementary inputs in the production of new ideas. The present model suggests, however, that this production process will be slowed by communication externalities when ideas are initially dispersed across people. This suggests a potentially novel way to think about the relationship between communication patterns and growth. An earlier version of this paper pursued this issue at greater length and obtained results on the relationship between the structure of communication networks and the rate of productivity growth.

4 Steady-State Learning

I next characterize the model's long-run behavior and its implications for technology adoption. Section 4.1 shows that the long-run probability distribution of knowledge endowments in the model is invariant within a large class of interaction structures of interest. Section 4.2 then establishes necessary and sufficient conditions for the learning process to converge to complete knowledge. One idea that emerges from this analysis is that certain kinds of complementary knowledge may serve as catalysts to promote the diffusion of other knowledge, and thus adoption; Section 4.3 illustrates this. Finally, Section 4.4 studies the mechanics of lock-in to one of several possible technologies.

4.1 Long-Run Equivalent Interaction Processes

As we turn to asymptotic behavior it will be important to understand how sensitive results are to the details of the interaction process. Any specification of the probability $q(t, \tau)$ that agent t speaks with agent $\tau < t$ generates a stochastic "communication tree". Figure 4 depicts an example: agent 2 spoke with agent 1, 3 with 2, 4 with 1, and so forth. In this Section the object of study will be the evolution of learning as paths within the communication tree elongate. Intuitively one expects that the evolution of the process $\{k_t\}$ will resemble the evolution of learning along a single infinitely-long branch provided that branches elongate "sufficiently quickly". This might not hold if, on the other hand, all agents $t > 1$ spoke with agent 1, in which case the model would be better described as a model of bilateral rather than social learning. One can rule this out for a large class of interaction processes that satisfy a monotonicity property:

Assumption 4. *$q(t, \tau)$ is non-decreasing in τ for every t.*

Recall that $q(t, \tau) = 0$ for $\tau \geq t$ since agents speak with predecessors; Assumption 4 says that they are at least weakly more likely to speak with more recent predecessors.

Figure 4: Stochastic Communication Tree at $t = 6$

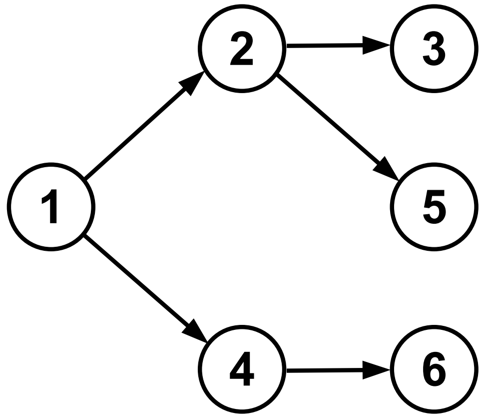

This property holds for most commonly-used interaction processes (e.g. random matching or death-and-replacement) and is sufficient to ensure that the long-run behavior of the process is like that of a strict Markov process in which $a(t) = t - 1$; hence it defines an equivalence class.

Proposition 4. *Let π_t be the distribution of k_t. There exists a unique π such that*

$$\lim_{t \to \infty} \pi_t = \pi \tag{3}$$

for any interaction structure $q(\cdot, \cdot)$ satisfying Assumption 4.

Proof. Proofs for this Section are deferred to Appendix A. □

The key to deriving this result is Lemma 5 in the Appendix, which shows that as t grows the probability that agent t is any finite number of conversations removed from agent 1 vanishes. I assume from here onwards that Assumption 4 holds and therefore that Proposition 4 applies.

4.2 When is Learning Asymptotically Complete?

This section characterizes situations in which long-run learning is complete, meaning that the probability that agent t has final knowledge endowment $\overline{k} \equiv (1, \ldots, 1)$ approaches 1 as $t \to \infty$. As in Section 3 the characterization is in terms of the value function V, treating the initial distribution F as a nuisance parameter drawn from the following class of admissible distributions:

Definition 2 (Admissibility)**.** *A distribution F on K is admissible if*

1. *For any b there exists $\tilde{k} \in support(F)$ such that $\tilde{k}^b = 1$*

2. $(0,\ldots,0) \in support(F)$

A sequence $\tilde{k}_1,\ldots$ *of initial endowments is admissible if there is an admissible distribution* F *such that each element of* $\tilde{k}_1,\ldots$ *occurs with positive probability under* F*.*

The first condition is innocuous; it essentially says that every unit of knowledge actually exists. The second requirement can be thought of as a robustness criterion – it focuses our attention on long-run behavior that is robust to the presence of agents with very little knowledge. For distributions F that do not satisfy this condition complete learning may obtain for a larger set of value functions V.

Given that short-run externalities involve too little communication, one expects that a necessary condition for complete long-run learning is that it be efficient. The next definition formalizes "efficient" for the long run:

Definition 3. *Complete learning is long-run efficient if, for any admissible sequence of initial endowments* $\tilde{k}_1,\ldots$ *such that* $\tilde{k}_1 \vee \ldots \vee \tilde{k}_\tau = \overline{k}$ *for some finite* τ*, supposing that* $a(t) = t-1$ *for all* $t > 1$*, then there exists* (t^*, T^*) *such that* $T > T^*$ *implies that any solution to*

$$\max_{\{k_t\}\ feasible} \sum_{t=1}^{T} V(k_t) - c\cdot(k_t - \tilde{k}_t)$$

has $k_t = \overline{k}$ *for* $t^* < t \leq T$*.*

The definition states that for any admissible sequence of initial knowledge endowments it will eventually be efficient for agents to learn everything.[14] The following lemma gives an equivalent characterization that is easier to read (and work with):

Lemma 1. *Complete learning is long-run efficient if and only if* $V(\overline{k}) - V(k) > c\cdot(\overline{k} - k)$ *for all* $k \prec \overline{k}$

Intuitively this says that it is always worth communicating enough to completely inform a recipient (though if V is supermodular it may not be worth communicating partial knowledge). The practical value of this characterization is that it abstracts entirely from the support of F, which facilitates the proof of this section's main proposition:

Proposition 5. *Long-run learning is complete under the following conditions:*

1. *If complete long-run learning is inefficient then it never obtains for any admissible* F*.*

[14] This definition is strong in the sense that there are sequences of knowledge endowments for which (a) the solution of the infinite-horizon problem with discounting would involve complete learning, while (b) there are many undiscounted T-horizon problems for which it does not. A countervailing benefit, however, is that we sidestep the vexed issue of discounting in OLG models. The requirement that *any* solution involve eventual complete learning is introduced to handle knife-edge cases in which tie-breaking is required; it parallels the assumption that agents break ties by choosing shorter messages.

2. *If complete long-run learning is efficient for a submodular V then it obtains for any admissible F.*

3. *If complete long-run learning is efficient for a supermodular V then it obtains for any admissible F if the following condition holds: there exists a sequence $k_1^*, \ldots, k_n^*$ satisfying $k_1^* = e_{b(1)}$, $k_n^* = \overline{k}$, and $k_t^* = k_{t-1}^* + e_{b(t)}$ for $t > 1$ such that $V(k_t^*) - V(m \vee e_{b(t)}) > c \cdot (k_{t-1}^* - m)$ for all $m \prec k_{t-1}^*$. If this condition does not hold then there exist admissible distributions F for which complete learning does not obtain.*

The first result is unsurprising, since agents in this model generally under-value communication compared to the social optimum. The second result is more meaningful but not particularly strong since (as the proof reveals) when complete learning is efficient for a submodular V there are in fact no externalities at all, even in the short run; in spite of decreasing returns, each unit of knowledge is always valuable enough to be worth communicating.

The results and discussion in Section 3.1 can be used to further characterize asymptotic welfare for submodular value functions when complete learning is not efficient. First consider the long-run behavior of the pest-control example in Section 3.1. Since the probability that a sequence of agents like that depicted in 2 arises is always positive, the limit average payoff in equilibrium must be bounded strictly below the limit average payoff a social planner would achieve. On the other hand, Proposition 2 implies that the limit average welfare loss is bounded above by $\sum c_b$.

Part 3 of Proposition 5 is of most interest. It says that decentralized communication may fail to assemble complementary units of knowledge unless a "small steps" condition holds. This condition is that we can order the n units of knowledge such that if an agent's initial knowledge endowment includes the $b(t)$'th unit of knowledge then it is strictly more efficient to communicate to the agent a message that results in him knowing all the $b(1), \ldots, b(t-1)$ previous units of knowledge than to communicate a message resulting in less information. This ensures that there is an order in which a knowledge endowment can be profitably augmented one unit of knowledge at a time. In addition to being sufficient for any F, this condition is necessary in situations where knowledge is initially highly dispersed. For example, if $V(1,1,1) = 4$, $V(k) = 0$ for all $k \neq (1,1,1)$, $c = (1,1,1)$, and agents are initially endowed with at most 1 unit of knowledge then the complementary units will never be assembled. Of course, in cases where some aggregation has already occurred so that the support of F includes larger elements then weaker conditions on V would be sufficient. Indeed, if $\overline{k} \in \text{support}(F)$ then no further conditions on V beyond the efficiency of complete learning would be necessary. In this case the issue facing society is less an aggregation problem than a search problem.[15]

[15]Situations where complete learning does eventually obtain may still involve large expected welfare losses if the rate of convergence is sufficiently slow. Not surprisingly the details of the interaction process $q(t, \tau)$

One implication is that larger-scale interactions may be particularly important when dealing with dispersed complementary knowledge. These could involve bringing more than two agents together at a time to communicate or simply increasing the number of bilateral interactions per agent. The "discovery" of James Watt's celebrated steam engine illustrates the latter. Watt began with Thomas Newcomen's design, which was operational but required too much coal to be economical. He then had two key social interactions. First, the chemist Joseph Black explained certain features of the process of evaporation and condensation to him, which inspired the addition of a separate condenser to the engine. Second, he acquired from the notorious John "Iron-mad" Wilkinson a new technique for smooth-boring shafts (originally developed to improve cannons) that greatly improved the fit of the pistons in Watt's cylinders. These complementary insights combined put Watt's engine over an efficiency threshold and made it a commercial success. Had Watt not spoken to either Black or Wilkinson, his name might be unknown today (Hart-Davis, 2001).

4.3 Catalytic Knowledge and Knowledge Policy

The "small steps" condition in Proposition 5 evokes the idea that some kinds of knowledge catalyze the diffusion of other complementary skills and thus may catalyze adoption. For example, general-purposes skills like numeracy complement many applied ones.

This idea can be brought out most clearly in an example with asymptotically incomplete learning. Again suppose that implementing IPM requires two kinds of knowledge, numeracy and the ability to identify insect eggs, but now suppose that while the latter can be communicated relatively easily, the former is prohibitively costly to teach. This is sensible, as most people spend many years in school acquiring numeracy, and consistent with Winarto (2004) who reports instances of social learning about eggs but not about mathematics. It implies that agents will adopt if and only if (1) they themselves are numerate, and (2) they or the peer with whom they speak knows about insect eggs. Crucially, the latter is more likely if the peer with whom they speak is himself numerate, since numerate people are more likely to acquire the complementary applied skill. This is consistent with evidence from diffusion research that more educated people are more likely to adopt innovations, that later adopters tend to learn from more educated peers, and that later adopters tend to learn from peers with more "change agent" contact (Rogers and Shoemaker, 1971, Generalizations 5-2, 6-3, and 6-6). The latter fact is readily explained by any reasonable model of social interaction, but the former suggest some form of filtering.

Under these assumptions long-run learning will be incomplete and there will be a

are key here; one can show by example (available on request) that expected losses may be either bounded or unbounded depending on q.

steady-state rate $\pi \in [0,1)$ of IPM adoption. Suppose that the distribution of initial knowledge endowments F has the following simple structure: a random fraction $s \in (0,1)$ of farmers get formal schooling and become numerate, while a random fraction e have contact with a government extension agent and learn about insect eggs. Then one can show that

$$\pi = es \underbrace{\left[1 + \frac{1-e}{1-(1-e)s}\right]}_{\text{Social Multiplier}} \tag{4}$$

The first term in the product is simply the probability that any one individual is endowed with the knowledge needed to perform IPM. The bracketed term is a social multiplier; it captures the way in which knowledge spillovers between farmers scale up the impact of investments in schooling s and extension e.

The interesting thing about this social multiplier is the way in which it depends asymmetrically on s and e. One can readily show that $\frac{\partial \pi}{\partial e} > 0$ and $\frac{\partial \pi}{\partial s} > 0$; unsurprisingly, both kinds of teaching increase adoption rates. This is consistent, for example, with the empirical evidence provided by Foster and Rosenzweig (1996), who show that Green Revolution technologies increased the return to schooling in India. The model also predicts decreasing returns to extension activity ($\frac{\partial^2 \pi}{\partial e^2} < 0$) because of a saturation effect: the more likely it is that a farmer has a neighbor who knows how to identify insect eggs, the less valuable it is for him to speak with an extension agent himself. However, there are *increasing* returns to schooling ($\frac{\partial^2 \pi}{\partial s^2} > 0$). Increasing returns result from the "weakest link" property of communication. As more farmers become numerate and value knowledge about insect eggs, that knowledge tends to travel further, and so the social cost incurred when its transmission halts because of an innumerate farmer increases. The socially optimal schooling policy will therefore be either to educate everyone or to educate no-one – in other words, it will tend towards egalitarianism.

If agents choose their education levels non-cooperatively then we have analogous results on incentives. The private returns to both schooling and extension contact are below the public returns, which potentially justifies subsidizing both. Individual investments in extension are strategic substitutes, but investments in schooling are strategic complements. If this complementarity is strong enough then there may exist multiple equilibria and a poverty trap in which society mis-coordinates on the low-schooling equilibrium. This suggests that general-purpose skills like numeracy may be doubly important for policy, both because of the incentive problems associated with providing them (Becker, 1993) and also because of their potential for large external returns.

Benhabib and Spiegel (2005) have argued that "the policy implications of distinguishing between the role of education as a factor of production and a factor that facilitates technology diffusion are significant" (p. 939). Our example illustrates this, and also shows that the two functions may be inseparable: numeracy facilities the diffusion if

IPM precisely because it complements IPM skills in the production function. This is a direct consequence of the fact that communication responds to incentives in the model – it would not hold if communication choices were random.

4.4 Multiple Steady States and Technology Lock-In

A final fundamental question about learning processes is whether or not they converge to a unique limit. This is particular important for understanding technology adoption, where lock-in to alternative technologies is often observed. As mentioned in the Introduction, Rogers and Kincaid (1981) report that women in each of 25 Korean villages tended to converge to use of the same method of birth control – the pill in one, IUDs in others, vasectomy in yet others. Can such patterns emerge due to filtering? While Proposition 4 implies that the ex-ante probability distribution π_t converges to a unique limit, *realizations* of the process $\{k_t\}$ may still converge to different limiting sets. This section (1) shows that this requires a non-supermodular value function and (2) provides an example of technological lock-in.

I use a notion of uniqueness adapted from the standard theory of Markov processes. Consider the special case where $a(t) = t - 1$ for all $t > 1$, so that $\{k_t\}$ is first-order Markov. Then it is well-known that the state space K can be uniquely partitioned as $K = T \cup L_1 \cup \ldots \cup L_J$ where (1) T is a set of transient states, meaning that $\mathbb{P}(k_t = k \text{ for some } t > 1 | k_1 = k) < 1$ for $k \in T$, and (2) each L_j has the properties (1) $k_t \in L_j$ implies $k_\tau \in L_j$ for all $\tau > t$ and (2) L_j contains no strict subset with this property. The process $\{k_t\}$ must eventually enter one of these sets L_j and will stay within it from then on; if there is exactly one such set then we can speak of the process having a unique limit.

To adapt this notion to more general non-Markovian processes I introduce the concept of agent t's descendants $D(t) = \{\tau : t \in A(\tau)\}$; these are the agents from whom a series of conversations trace back to t. Call a subset $L \subseteq K$ a limit set if $k_t \in L$ implies $k_\tau \in L$ for all $\tau \in D(t)$ and L does not contain any proper subsets with this property. It is no longer the case that $k_t \in L_j \Rightarrow k_{t'} \in L_j$ for all $t' > t$, since it is possible that agent t' speaks with a predecessor t'' with $k_{t''} \notin L_j$. However, if the chain has a unique limit set L then Proposition 4 implies that $\lim_{t\to\infty} \mathbb{P}(k_t \in L) = 1$, so that we can speak of L as the unique limit of the process regardless of the details of the interaction structure q. Our next result shows that this is the case whenever V is supermodular.

Proposition 6. *If V is supermodular then there exists a limit set $L \subseteq K$ such that $\lim_{t\to\infty} \mathbb{P}(k_t \in L) \to 1$.*

This may at first seem somewhat surprising given the association between complementarity and multiple equilibria in other models, but the essential idea is simple: for there to be path-dependence it must hold that the accumulation of one kind of knowledge tends to discourage the accumulation of another kind, while under complementarity accumulation of any one kind of knowledge weakly *encourages* the accumulation of other kinds. Path-dependence thus requires some degree of substitutability.

Learning about alternative technologies has exactly this feature. Suppose that a worker can choose between two software packages, A and B, and that his productivity using either package depends on how skilled he is using it. Let the n units of knowledge be skills specific to these packages, with skills $1, \dots, n/2$ specific to alternative A and the rest specific to B. A worker's productivity using a particular package is proportional to the number of specific skills he knows, with potentially different proportionality constants α_A and α_B capturing intrinsic productivity differences. The worker's value function is

$$V(k) = \max \left\{ \alpha_A \sum_{b=1}^{n/2} k^b, \alpha_B \sum_{b=n/2+1}^{n} k^b \right\} \tag{5}$$

Communicating each skill costs c and $0 < c < \alpha_z$ for $z \in \{A, B\}$ so that it is locally optimal to communicate a skill if and only if the recipient will use it. Workers' initial skills are concentrated in one or the other technology (this is privately optimal and also broadly consistent with the outcome of social learning dynamics, as we shall see). Each worker t knows a random subset of the skills relevant for one of the technologies, and let l_z be the *most* that any worker initially knows about technology $z \in \{A, B\}$.

Are there cases in which either technology could be selected with positive probability? There is always positive probability of a finite sequence $\tilde{k}_1, \dots, \tilde{k}_T$ occurring such that agent T knows *all* the skills for either one of the technologies, so the key question becomes: under what conditions will such an agent still choose with positive probability to teach a peer about the other technology? The net local payoff from teaching about technology z is at least $\frac{n}{2}(\alpha_z - c)$ (if the learner knows nothing about z). As for teaching z', if we start in a steady state in which agents have been communicating exclusively about z then the teacher will be able to teach at most $l_{z'}$ things about z', and so the net local payoff from doing this is at most

$$\alpha_{z'} l_{z'} + (\alpha_{z'} - c) \min \left\{ \frac{n}{2} - l_{z'}, l_{z'} \right\} \tag{6}$$

The first term is the payoff from skills the learner was endowed with, and the second the payoff from skills transferred to him. We conclude that adoption of z is an absorbing state if

$$\frac{n}{2}(\alpha_z - c) > \alpha_{z'} l_{z'} + (\alpha_{z'} - c) \min \left\{ \frac{n}{2} - l_{z'}, l_{z'} \right\} \tag{7}$$

As one would expect, this is more likely to hold when technology z is innately more productive relative to z' (i.e. α_z large relative to $\alpha_{z'}$). Interestingly, it is also more likely to hold as the total amount of available knowledge $n/2$ increases above the amount $l_{z'}$ that any one person initially knows. Either technology outcome becomes more stable as knowledge is initially more disaggregated, because the likelihood of any one person knowing enough to warrant a technology switch shrinks. Finally, lock-in on z becomes (weakly) less likely as the cost c of communication increases, because this lowers the benefits of teaching an agent accumulated skills relative to letting him use what he already knows.

One can verify that (7) may hold for both $z = A$ and $z = B$, in which case the process has multiple absorbing states. In particular, many or all agents may become locked in to an inferior technology. Further examination of (7) shows that as n grows large the productivity gap $\alpha_{z'} - \alpha_z$ can grow arbitrarily large and yet adoption of z remains an equilibrium. The problem is that once one technology becomes sufficiently well-understood it is no longer privately optimal to communicate about the other one, even if enough discussion about the other could lead to its re-adoption. Consequently there may be situations in which society would be better-off with either more or less communication amongst early movers. This holds even though in some sense all communication externalities are driven by under-communication; here, under-communication of knowledge about one technology is driven by earlier communication about the alternative.

Set against the literature on adoption, the interesting feature of this example is that lock-in occurs even though there is no underlying payoff interdependency and no uncertainty about the returns to either technology. Arthur (1989), Kremer and Miguel (2007), and Brock and Durlauf (2010) are examples of models in which one agent's returns from adopting depend on other agents' adoption choices, potentially leading to multiple equilibria. Uncertainty can lead to lock-in through herding (Banerjee, 1992; Bikhchandani et al., 1992), through the use of decision rules that favor imitation (Ellison and Fudenberg, 1993, 1995), or through the standard bandit mechanism (Bala and Goyal, 1998). The filtering model lacks both features but still yields lock-in, suggesting it may be a more general phenomenon than previously appreciated.

5 Conclusion

For standard goods and services, functioning markets help to ensure that localized optimization leads to aggregate efficiency. Markets for knowledge area rarely feasible, however; instead much of our knowledge flows to us through social interactions. A key issue is under what conditions individual choices about what knowledge to communicate aggregate up to efficient social learning. This paper addresses these question using a new analytic framework which explicitly accounts for communication costs and characterizes their implications for efficiency and for outcomes like technology adoption.

The perspective adopted here, that communication responds to local costs and benefits, could be fruitfully applied to many other issues. One example is the design of messages such as product advertisements or social marketing campaigns. A theory of filtering permits a distinction between messages that are *individually* persuasive and those that are likely to "go viral" and become *socially* persuasive as well. Organizations are another important area of application. Recent models have assumed contractible communication and studied how to design organizations around real constraints (Bolton and Dewatripont, 1994; Garicano, 2000; Crémer et al., 2007); locally optimal communication may provide a way to understand and model incentive constraints.

Finally, while this paper has focused on learning *how* to do things, filtering is also important for learning *whether* to do things – whether or not to buy a product, invest in education, or adopt a technology when the returns are uncertain. The intuitions developed here should transfer to such settings. For example, a technology with many small benefits may take longer to catch on than a technology with one "killer app," because knowledge of these benefits is complementary.

References

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar, "Communication Dynamics in Endogenous Social Networks," Technical Report, Massachusetts Institute of Technology 2010.

—, **Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar**, "Bayesian Learning in Social Networks," NBER Working Papers 14040, National Bureau of Economic Research, Inc May 2008.

Ambrus, Attila, Eduardo Azevedo, and Yuichiro Kamada, "Hierarchical Cheap Talk," Technical Report, Harvard University May 2010.

Arrow, Kenneth, "Economic Welfare and the Allocation of Resources for Invention," in R. Nelson, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 1962, pp. 609–626.

Arthur, W Brian, "Competing Technologies, Increasing Returns, and Lock-In by Historical Events," *Economic Journal*, March 1989, *99* (394), 116–31.

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

Banerjee, Abhijit and Drew Fudenberg, "Word-of-mouth learning," *Games and Economic Behavior*, January 2004, *46* (1), 1–22.

Banerjee, Abhijit V, "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, August 1992, *107* (3), 797–817.

—, "The Economics of Rumours," *Review of Economic Studies*, April 1993, *60* (2), 309–27.

Becker, Gary, *Human Capital: a Theoretical and Empirical Analysis with Special Reference to Education*, Chicago: University of Chicago Press, 1993.

Benhabib, Jess and Mark M. Spiegel, "Human Capital and Technology Diffusion," in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Vol. 1 of *Handbook of Economic Growth*, Elsevier, 2005, chapter 13, pp. 935–966.

Bentley, Jeffrey, "Folk experiments," *Agriculture and Human Values*, 2006, *23*, 451–462.

Besley, Timothy and Anne Case, "Diffusion as a Learning Process: Evidence from HYV Cotton," Working Papers 228, Princeton University, Woodrow Wilson School of Public and International Affairs, Research Program in Development Studies. May 1994.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades," *Journal of Political Economy*, October 1992, *100* (5), 992–1026.

Black, Duncan and Vernon Henderson, "A Theory of Urban Growth," *Journal of Political Economy*, April 1999, *107* (2), 252–284.

Bollobás, Belá, *Random Graphs*, 2nd. ed., Cambridge University Press, 2001.

Bolton, Patrick and Christopher Harris, "Strategic Experimentation," *Econometrica*, March 1999, *67* (2), 349–374.

_ **and Mathias Dewatripont**, "The Firm as a Communication Network," *The Quarterly Journal of Economics*, November 1994, *109* (4), 809–39.

Bramoulle, Yann and Rachel Kranton, "Public goods in networks," *Journal of Economic Theory*, July 2007, *135* (1), 478–494.

Brock, William A. and Steven N. Durlauf, "Adoption Curves and Social Interactions," *Journal of the European Economic Association*, 03 2010, *8* (1), 232–251.

Calvó-Armengol, Antoni and Joan de Martí, "Communication Networks: Knowledge and Decisions," *American Economic Review*, May 2007, *97* (2), 86–91.

Conley, Timothy G. and Christopher Udry, "Learning about a New Technology: Pineapple in Ghana," *American Economic Review*, 2010, *100* (1), 35–69.

Crémer, Jacques, Luis Garicano, and Andrea Prat, "Language and the Theory of the Firm," *The Quarterly Journal of Economics*, 02 2007, *122* (1), 373–407.

DeGroot, Morris, "Reaching a Consensus," *Journal of the American Statistical Association*, March 1974, *69* (345), 118–121.

Demarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion Bias, Social Influence, And Unidimensional Opinions," *The Quarterly Journal of Economics*, August 2003, *118* (3), 909–968.

Dewatripont, Mathias and Jean Tirole, "Modes of Communication," *Journal of Political Economy*, December 2005, *113* (6), 1217–1238.

Duflo, Esther and Emmanuel Saez, "The Role Of Information And Social Interactions In Retirement Plan Decisions: Evidence From A Randomized Experiment," *The Quarterly Journal of Economics*, August 2003, *118* (3), 815–842.

Ellison, Glenn and Drew Fudenberg, "Rules of Thumb for Social Learning," *Journal of Political Economy*, August 1993, *101* (4), 612–43.

_ and _ , "Word-of-Mouth Communication and Social Learning," *The Quarterly Journal of Economics*, February 1995, *110* (1), 93–125.

Feldman, Maryann P, "Knowledge Complementarity and Innovation," *Small Business Economics*, October 1994, *6* (5), 363–72.

Foster, Andrew D and Mark R Rosenzweig, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, December 1995, *103* (6), 1176–1209.

_ and _ , "Technical Change and Human-Capital Returns and Investments: Evidence from the Green Revolution," *American Economic Review*, September 1996, *86* (4), 931–53.

Garicano, Luis, "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, October 2000, *108* (5), 874–904.

Glaeser, Edward L., "Learning in Cities," *Journal of Urban Economics*, September 1999, *46* (2), 254–277.

_ and Cass R. Sunstein, "Extremism and Social Learning," NBER Working Papers, National Bureau of Economic Research, Inc December 2007.

Gramlich, Edward M, "Infrastructure Investment: A Review Essay," *Journal of Economic Literature*, September 1994, *32* (3), 1176–96.

Grimmett, Geoffrey and David Stirzaker, *Probability and Random Processes*, New York: Oxford University Press, 2002.

Hart-Davis, Adam, "James Watt and the Lunaticks of Birmingham," *Science*, April 2001, *292* (5514), 55–56.

Heckman, James and Pedro Carneiro, "Human Capital Policy," NBER Working Papers, National Bureau of Economic Research, Inc February 2003.

Holton, Gerald, *The Scientific Imagination*, Cambridge, MA: Harvard University Press, 1998.

Jackson, Matthew O. and Benjamin Golub, "Nave Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, February 2010, *2* (1), 112–149.

Jacobs, Jane, *The Economy of Cities*, New York, NY: Random House, 1969.

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson, "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *The Quarterly Journal of Economics*, August 1993, *108* (3), 577–98.

Jensen, Robert, "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector," *The Quarterly Journal of Economics*, 08 2007, *122* (3), 879–924.

Jovanovic, Boyan and Rafael Rob, "The Growth and Diffusion of Knowledge," *Review of Economic Studies*, October 1989, *56* (4), 569–82.

_ **and Yaw Nyarko**, "The transfer of human capital," *Journal of Economic Dynamics and Control*, 1995, *19* (5-7), 1033–1064.

Kondor, Peter and Gergely Ujhelyi, "Social Learning with Subjective Communication and Self-Selection," May 2005. mimeo.

Kremer, Michael and Edward Miguel, "The Illusion of Sustainability," *The Quarterly Journal of Economics*, 08 2007, *122* (3), 1007–1065.

Lucas, Robert Jr., "On the mechanics of economic development," *Journal of Monetary Economics*, July 1988, *22* (1), 3–42.

Mankiw, N. Gregory and Michael D. Whinston, "Free Entry and Social Inefficiency," *RAND Journal of Economics*, Spring 1986, *17* (1), 48–58.

Marmaros, David and Bruce Sacerdote, "How Do Friendships Form?," *The Quarterly Journal of Economics*, 02 2006, *121* (1), 79–119.

Milgrom, Paul and Chris Shannon, "Monotone Comparative Statics," *Econometrica*, January 1994, *62* (1), 157–80.

Mokyr, Joel, "Punctuated Equilibria and Technological Progress," *American Economic Review*, May 1990, *80* (2), 350–54.

Moretti, Enrico, "Human capital externalities in cities," in J. V. Henderson and J. F. Thisse, eds., *Handbook of Regional and Urban Economics*, Vol. 4 of *Handbook of Regional and Urban Economics*, Elsevier, 06 2004, chapter 51, pp. 2243–2291.

— , "Social Learning and Peer Effects in Consumption: Evidence from Movie Sales," *Review of Economic Studies*, 2011, *78* (1), 356–393.

Munshi, Kaivan, "Social learning in a heterogeneous population: technology diffusion in the Indian Green Revolution," *Journal of Development Economics*, February 2004, *73* (1), 185–213.

Nelson, Richard R. and Edmund S. Phelps, "Investment in Humans, Technological Diffusion, and Economic Growth," *The American Economic Review*, 1966, *56* (1/2), pp. 69–75.

Rogers, Everett, *Diffusion of Innovations*, 5th ed. 2003.

— **and D. Lawrence Kincaid**, *Communication Networks: Towards a New Paradigm for Research*, New York: Free Press, 1981.

— **and Floyd Shoemaker**, *Communication of Innovations: A Cross-Cultural Approach*, 2nd. ed. 1971.

Romer, Paul M, "Increasing Returns and Long-run Growth," *Journal of Political Economy*, October 1986, *94* (5), 1002–37.

Sinha, B.P. and Prayag Mehta, "Farmers' Need for Achievement and Change-Proneness in Acquisition of Information from a Farm-Telecast," *Rural Sociology*, 1972, *37* (3).

Smith, Lones and Peter Sorensen, "Pathological Outcomes of Observational Learning," *Econometrica*, March 2000, *68* (2), 371–398.

Stein, Jeremy, "Conversations among Competitors," *American Economic Review*, December 2008, *98(5)*, 215062.

Straub, Stephane, "Infrastructure and development : a critical appraisal of the macro level literature," Policy Research Working Paper Series 4590, The World Bank April 2008.

Topkis, Donald M., "Minimizing a Submodular Function on a Lattice," *Operations Research*, March-April 1978, *26* (2), 305–321.

Weitzman, Martin L., "Recombinant Growth," *The Quarterly Journal of Economics*, May 1998, *113* (2), 331–360.

Winarto, Yunita, *Seeds of Knowledge: The Beginning of Integrated Pest Management in Java*, New Haven: Yale University Southeast Asian Studies, 2004.

A Proofs

Theorems 3.1 and 3.2 of Topkis (1978) establish that V is supermodular (submodular) if and only if it exhibits increasing (decreasing) differences; I generally use the latter formulation.

A.1 Useful Lemmas

Lemma 2. *If V is supermodular then $C(k_s, \tilde{k}_r)$ is unique.*

Proof. Suppose towards a contradiction that both k', k'' satisfy

$$k \in \arg\max_{k \in K_r} V(k) - c \cdot (k - \tilde{k}_r)$$

where as before $K_r = \{k : \tilde{k}_r \precsim k \precsim \tilde{k}_r \vee k_s\}$. Let $d'' = k'' \setminus k'$; k' and k'' must be un-orderable so $|d''| > 0$. By local optimality (and invoking the tie-breaking assumption)

$$V(\tilde{k}_r \vee (k'' \setminus d'') \vee d'') - V(\tilde{k}_r \vee (k'' \setminus d'')) > c \cdot d''$$

By supermodularity

$$V(\tilde{k}_r \vee k' \vee d'') - V(\tilde{k}_r \vee k') \geq V(\tilde{k}_r \vee (k'' \setminus d'') \vee d'') - V(\tilde{k}_r \vee (k'' \setminus d''))$$

since by definition $k' \succsim k'' \setminus d''$. But combining these equations yields $V(\tilde{k}_r \vee k' \vee d'') - V(\tilde{k}_r \vee k') > c \cdot d''$, contradicting the local optimality of k'. □

Lemma 3. *If V is supermodular and $k'_s \succsim k_s$ then $C(k'_s, \tilde{k}_r) \succsim C(k_s, \tilde{k}_r)$.*

Proof. By Lemma 2 C is uniquely defined by

$$C(k_s, \tilde{k}_r) = \arg\max_{k \in K_r(k_s)} V(k) - c \cdot (k - \tilde{k}_r) \tag{8}$$

where $K_r(k_s) = \{k : \tilde{k}_r \precsim k \precsim \tilde{k}_r \vee k_s\}$. The objective function is supermodular since V is supermodular and costs are linear; by Theorem 4 of Milgrom and Shannon (1994) it is sufficient to show that the constraint sets $K_r(k_s)$ are increasing in k_s in the strong set order.

Suppose therefore that $k'_s \succsim k_s$ and let $k \in K_r(k_s)$, $k' \in K_r(k'_s)$. We know $k \wedge k' \precsim k \precsim k_s \vee \tilde{k}_r$. Since both $k \succsim \tilde{k}_r$ and $k' \succsim \tilde{k}_r$ we must have $k \wedge k' \succsim \tilde{k}_r$. Together these facts establish $k \wedge k' \in K_r(k_s)$. We know $k \vee k' \succsim k \succsim \tilde{k}_r$. Since both $k' \precsim k'_s \vee \tilde{k}_r$ and $k \precsim k_s \vee \tilde{k}_r \precsim k'_s \vee \tilde{k}_r$ we must have $k \vee k' \precsim k'_s \vee \tilde{k}_r$. Together these facts establish $k \vee k' \in K_r(k'_s)$. This establishes that $K_r(\cdot)$ is increasing in the strong set order. □

Lemma 4. *Consider a sequence* $\{z^r\}$ *on the* n*-dimensional simplex with* $\lim_{r\to\infty} z^r = z$ *and a sequence* $\{y^t\}$ *such that* y^t *is on the* $t-1$*-dimensional simplex, with* y^t_r *denoting the* r*th component. The series* $x_t = \sum_{r<t} y^t_r z^r$ *evolves on the* n*-dimensional simplex. Let* $Y^t_r = \sum_{s\le r} y^t_s$*. If* $\lim_{t\to\infty} Y^t_r = 0$ *for every finite* r*, then* $\lim_{t\to\infty} x_t = z$*.*

Proof. Fix $\epsilon > 0$ and decompose

$$|x_t - z| = \left|\sum_{s\le r} y^t_s(z^s - z) + \sum_{s=r+1}^{t} y^t_s(z^s - z)\right|$$

Since $z^r \to z$, $\exists \underline{r}$ such that $r > \underline{r}$ implies $|z^r - z| \le \frac{\epsilon}{2}$. Moreover since $Y^t_{\underline{r}} \to 0$ for this $\underline{r}$, $\exists \underline{t}(\underline{r}) > \underline{r}$ such that $t > \underline{t}(\underline{r})$ implies $Y^t_{\underline{r}} \le \frac{\epsilon}{2}$. For $t > \underline{t}(\underline{r})$,

$$|x_t - z| \le \sum_{s\le \underline{r}} y^t_s\,|z^s - z| + \sum_{s=\underline{r}+1}^{t} y^t_s\,|z^s - z|$$

by the triangle inequality. $|z^s - z| \le 1$ so the left-hand term is at most $Y^t_{\underline{r}} \le \frac{\epsilon}{2}$; $|z^s - z| \le \frac{\epsilon}{2}$ for $s > \underline{r}$, so the right-hand term is at most $\frac{\epsilon}{2}$. Thus $|x_t - z| \le \epsilon$. □

Lemma 5. *Let* p^t_d *be the probability that agent* t *is* d *steps from agent 1 in the communication tree, and let* $P^t_d = \sum_{s\le d} p^t_s$ *be the probability that agent* t *is at most* d *steps from agent 1. Then for any fixed* d*,* $\lim_{t\to\infty} P^t_d = 0$*.*

Proof. Agent t is matched to some agent $\tau < t$ according to the distribution $q(t,\tau)$. The a priori probability p^t_d that agent t is d steps away from agent 1 is

$$p^t_d = \sum_{r<t} q(t,r)p^r_{d-1}$$

and so the probability of being at most d steps away from agent 1 is

$$\begin{aligned}
P^t_d &= \sum_{s\le d}\sum_{r<t} q(t,r)p^r_{s-1} \\
&= \sum_{r<t}\sum_{s\le d} q(t,r)p^r_{s-1} \\
&= \sum_{r<t} q(t,r)\sum_{s\le d} p^r_{s-1} \\
&= \sum_{r<t} q(t,r)\sum_{s\le d-1} p^r_{s} \\
&= \sum_{r<t} q(t,r)P^r_{d-1}
\end{aligned}$$

It is trivial that $P^t_0 = 0$ for any agent $t > 1$ (all agents must be at least one hop away from the root node) and therefore $\lim_{t\to\infty} P^t_0 = 0$. We will show by induction that P^t_u must also approach 0 for any $u > 0$. Suppose then that for $u \leq d-1$ we know $\lim_{t\to\infty} P_u^t = 0$. Then by applying Lemma 4 to the last line above, it will also hold that $\lim_{t\to\infty} P_d^t = 0$ provided that

$$\lim_{t\to\infty} \sum_{s\leq\tau} q(t,s) = 0$$

Since $q(t,\tau)$ is weakly increasing in τ we must have

$$\sum_{s\leq\tau} q(t,s) \leq \frac{\tau}{t}$$

which carried into the limit gives the desired result. Therefore by induction we have $\lim_{t\to\infty} P_d^t = 0$ for any finite d as desired. □

A.2 Proof of Proposition 4

Proof. The approach is as follows: first, show that π_t converges to a limit π when q has the first-order Markov structure

$$q(t,\tau) = \begin{cases} 1 & \tau = t-1 \\ 0 & \tau \neq t-1 \end{cases} \tag{9}$$

and then show that if this holds then π_t must converge to the same π for any other admissible q.

First, let q be as above so that $\{k_t\}$ is a first-order Markov chain. It is well-known (e.g. Theorem 6.4.21 of Grimmett and Stirzaker (2002)) that the corresponding probabilities $\pi_t(k)$ converge if the chain is aperiodic, and since π_t is finite dimensional this also implies convergence of π_t using the Euclidean metric. To establish aperiodicity, note that if $\tilde{k}_{t+1} = \tilde{k}_t$ then by local optimality we must have $k_{t+1} = k_t$. Consequently for any state k that is reached with positive probability it is possible to return to k after m periods for any integer m. Ignoring states that are never reached, this implies that the entire process $\{k_t\}$ is aperiodic as required.

Now consider an arbitrary admissible q and let π_t be the a priori distribution over K for agent t. Let $\pi(d)$ be the corresponding a priori probability distribution for an agent *conditional* on being d steps removed from the root agent 1; by appeal to the above result for Markov chains, $\pi(d) \to \pi$ for some π as $d \to \infty$. Finally, let p_d^t be the probability that agent t is d steps from the root agent. Then

$$\pi_t = \sum_{\tau<t} p_d^\tau \pi(d) \tag{10}$$

By Lemma 5, $\lim_{t\to\infty} P_d^t = 0$. By Lemma 4, this along with $\pi(d) \to \pi$ as $d \to \infty$ implies

$\lim_{t\to\infty} \pi_t = \pi$ as well. $\square$

A.3 Proof of Lemma 1

If. Consider any admissible sequence of knowledge endowments and let τ be such that $\tilde{k}_1 \vee \ldots \vee \tilde{k}_\tau = \overline{k}$. $V(\overline{k}) - V(k) > c \cdot (\overline{k} - k)$ clearly implies that if $k_t = \overline{k}$ then it must be efficient to have $k_{t'} = \overline{k}$ for all $t' > t$. Thus in order for complete learning to be long-run *inefficient* it would have to be that the optimum involved $k_t \prec \overline{k}$ for all $\tau < t \leq T$ and for arbitrarily large T. The welfare change from switching from this policy to one in which all agents communicate everything to their successors is at least

$$\left[\sum_{t=\tau+1}^{T} \left(V(\overline{k}) - V(k_t) - c \cdot (\overline{k} - k_t)\right)\right] - \left[\tau \sum_b c^b\right] \tag{11}$$

where the summation captures the benefits to agents indexed greater than τ and the last term is an upper bound on the costs to agents indexed less than τ. Letting $B = \min_{k \prec \overline{k}} V(\overline{k}) - V(k) - c \cdot (\overline{k} - k)$, we obtain a lower bound on the welfare change equal to $B(T - \tau) - \tau \sum_b c^b$, which is positive for sufficiently large T.

Only If. Suppose there exists some $k \prec \overline{k}$ such that $V(\overline{k}) - V(k) \leq c \cdot (\overline{k} - k)$. For any T, $\tilde{k}_T \precsim k$ occurs with positive probability (since $(0, \ldots, 0) \in \text{support}(F)$). If $V(\overline{k}) - V(k) < c \cdot (\overline{k} - k)$ then $k_T = \overline{k}$ is inefficient; if $V(\overline{k}) - V(k) = c \cdot (\overline{k} - k)$ then there is an efficient communication pattern with $k_T \prec \overline{k}$. Either case contradicts the definition of long-run efficient learning.

A.4 Proof of Proposition 5

Long-run Learning is Inefficient. By Lemma 1 this implies there exists k such that $V(\overline{k}) - V(k_t) < c \cdot (\overline{k} - k)$. For any t, $\tilde{k}_t \precsim k$ with positive probability, since $(0, \ldots, 0) \in \text{support}(F)$. Regardless of whom t speaks with, $k_t \prec \overline{k}$ by local optimality.

Long-run Learning is Efficient for Submodular V. I will show that the unique limit on a chain is complete learning and then appeal to Proposition 4 to extend the result to general processes satisfying Assumption 4. Fix any k_t and any $\tilde{k}_{t+1}$. Suppose towards a contradiction that $k_{t+1} \prec k_t \vee \tilde{k}_{t+1}$ and let $d \equiv \tilde{k}_{t+1} \vee k_t - k_{t+1}$ be the knowledge not communicated to $t+1$. By submodularity, $V(k_{t+1} \vee d) - V(k_{t+1}) \geq V(\overline{k}) - V(\overline{k} \setminus d)$. By Lemma 2, $V(\overline{k}) - V(\overline{k} \setminus d) \geq c \cdot d$. Together these imply a contradiction of local optimality. Thus $k_{t+1} = k_t \vee \tilde{k}_{t+1}$ and complete learning eventually obtains with probability 1.

Long-run Learning is Efficient for Supermodular V. Proposition 4 establishes that there is a unique long-run limit and that it can be found by studying the linear case in which agent t speaks with agent $t-1$. It remains to establish the necessary and sufficient conditions for this limit to be complete learning.

If. Suppose there exists a sequence $\{k_t^*\}$ with the given properties. By admissibility of F, $k_1 = \tilde{k}_1 \succsim k_1^*$ occurs with positive probability. Inductively suppose that the event $\{k_t \succsim k_t^*\}$ occurs with positive probability for some t. Suppose $k_{t+1} \succsim k_{t+1}^*$ does not hold; then the message m from t to $t+1$ must satisfy $m \prec k_t^*$, and if $d = k_t^* - m$ it must be that $|d| > 0$. Then

$$\begin{aligned} V(k_{t+1} \vee d) - V(k_{t+1}) - c \cdot d &\geq V(k_{t+1}^*) - V(k_{t+1}^* \setminus d) - c \cdot d \\ &= V(k_{t+1}^*) - V((k_t^* - d) \vee e_{b(t+1)}) - c \cdot d \\ &> 0 \end{aligned}$$

The first weak inequality follows from supermodularity and $k_{t+1} \succsim k_{t+1}^* \setminus d$; the second strict inequality inequality is by assumption, since $|d| > 0$. The result contradicts local optimality; we can conclude that $k_{t+1} \succsim k_{t+1}^*$. Thus by induction an entire sequence $\{k_t\}$ such that $k_t \succsim k_t^*$ arises with positive probability, and in particular $\overline{k}$ is reached with positive probability. Relevance requires that $V(\overline{k}) - V(k) > c \cdot (\overline{k} - k)$ for all k so that once reached the state $\overline{k}$ persists and thus $\{\overline{k}\}$ constitutes a limit set. Finally by Proposition 6 this limit is unique.

Only If. It is enough to show that the given conditions are necessary for some particular class of admissible distributions F; here I consider those such that $|k| = 1$ for all $k \in \text{support}(F)$. Suppose that in the unique limit $k_t = \overline{k}$. Then there must exist a finite sequence $\tilde{k}_1, \ldots, \tilde{k}_T$ generating $k_1, \ldots, k_T$ such that $k_T = \overline{k}$. The proof proceeds by repeatedly pruning this sequence to obtain a subsequence with the stated characteristics.

First, prune elements from the beginning of the chain until $k_2 = \tilde{k}_1 + \tilde{k}_2$. Next, suppose that there exists $1 < \tau < T$ such that $m_{\tau+1} \equiv k_{\tau+1} - \tilde{k}_{\tau+1} \precsim k_\tau - \tilde{k}_\tau \equiv m_\tau$. m_τ denotes the "message" or knowledge passed in conversation to agent τ, and we have defined a situation where knowledge passed to τ is *not* passed to $\tau + 1$. Consider the modified sequence $\tilde{k}_1, \ldots, \tilde{k}_{\tau-1}, \tilde{k}_{\tau+1}, \ldots, \tilde{k}_T$ generating $\hat{k}_1, \ldots, \hat{k}_{\tau-1}, \hat{k}_{\tau+1}, \ldots, \hat{k}_T$. I claim that $\hat{k}_{\tau+1} \succsim k_{\tau+1}$. Let $d = k_{\tau+1} \setminus \hat{k}_{\tau+1} \precsim m_{t+1}$ and suppose towards a contradiction that $|d| > 0$. Then

$$\begin{aligned} V(\hat{k}_{\tau+1} \vee d) - V(\hat{k}_{\tau+1}) &= V(\hat{k}_{\tau+1} \vee (k_{\tau+1} \setminus d) \vee d) - V(\hat{k}_{\tau+1} \vee (k_{\tau+1} \setminus d)) \\ &\geq V((k_{\tau+1} \setminus d) \vee d) - V(k_{\tau+1} \setminus d) \\ &> c \cdot d \end{aligned}$$

where the first inequality follows from supermodularity and the second from local optimality in the original sequence. But since $d \precsim m_{t+1} \precsim m_t \precsim k_{\tau-1} = \hat{k}_{\tau-1}$ it is feasible to send d to agent $\tau + 1$ in the modified sequence, so we have a violation of local optimality in the modified sequence. Hence $|d| = 0$, i.e. $\hat{k}_{\tau+1} \succsim k_{\tau+1}$. By repeated application of

Lemma 3 it must then hold that the modified sequence still finishes with $\hat{k}_T = \overline{k}$.

Finite repetition of this pruning yields a chain $\tilde{k}'_1, \ldots, \tilde{k}'_{T'}$ in which $m'_t \succ m'_{t-1}$ for all $1 < t \leq T'$ – or in other words, $m'_t = k'_{t-1}$ and $k'_t = k'_{t-1} + \tilde{k}'_t$ – and still $k'_{T'} = \overline{k}$. By local optimality $V(k'_t) - V(k) > c \cdot (k'_t - k)$ for any $k \prec m'_t = k'_{t-1}$. The sequence $\{k'_t\}$ thus satisfies the desired conditions.

A.5 Proof of Proposition 6

Proof. By Proposition 4 it is sufficient to prove the result for the case in which agent t speaks with agent $t-1$, so that $\{k_t\}$ is a Markov chain on a finite state space K. It is well-known (see for example Grimmett and Stirzaker (2002)) that it eventually reaches a limit set $L \subseteq K$ such that all elements in L intercommunicate and no element in L communicates with an element not in L. Suppose towards a contradiction that two such limit sets L and L' exist. Let l and l' be maximally-valued elements within them, i.e. let $V(l) \geq V(\hat{l})$ for all $\hat{l} \in L$, and similarly for l'. If both limit sets are reached with positive probability then there exist endowment sequences $\{\tilde{k}_t\}_{t=1}^{T}$ and $\{\tilde{k}'_t\}_{t=1}^{T'}$ generating knowledge sequences $\{k_t\}_{t=1}^{T}$ and $\{k'_t\}_{t=1}^{T'}$ that reach l and l', respectively. Without loss of generality let $V(l) \leq V(l')$.

Now consider the evolution of the process beginning at $k_\tau = l$, and suppose that that (as occurs with positive probability) ensuing draws from the support of F yield the sequence $\{\tilde{k}'_t\}$. Trivially $k_{\tau+1} \succsim \tilde{k}'_1$, and by repeated application of Lemma 3 $k_{\tau+T'} \succsim l'$, which implies $V(k_{\tau+T'}) \geq V(l')$. If this is strict then $V(k_{\tau+T'}) > V(l') \geq V(l)$ which contradicts the maximality of $V(l)$ on L; if instead $V(k_{\tau+T'}) = V(l')$. If $k_{\tau+T'} \succ l'$ then since $l' \succsim \tilde{k}'_{T'}$ the difference $k_{\tau+T'} - l'$ must be contained in $k_{\tau+T'-1}$, but then the conversation between $\tau + T' - 1$ and $\tau + T'$ violated local optimality. So $k_{\tau+T'} = l'$, and therefore l communicates with l', contradicting the definition of L. □