The Internet of Things and Information Fusion: Who Talks to Who?

Soroush Saghafian[1], Brian Tomlin[2], Stephan Biller[3]
[1]Harvard Kennedy School, Harvard University, Cambridge, MA
[2]Tuck School of Business at Dartmouth, Hanover, NH
[3]CEO, Advanced Manufacturing International, Clearwater, FL

Manufacturing and Service Operations Management (M&SOM), Forthcoming

Problem definition: Autonomous sensors connected through the Internet of Things (IoT) are deployed by different firms in the same environment. The sensors measure an important operating-condition state variable but their measurements are noisy and so estimates are imperfect. Sensors can improve their own estimates by soliciting estimates from other sensors. The choice of which sensors to communicate with ("target") is challenging because sensors (a) are constrained in the number of sensors they can target, and (b) only have partial knowledge of how other sensors operate, i.e., they do not know others' underlying inference algorithms/ models. We study the targeting problem, examine the evolution of inter-firm sensor communication patterns, and explore what drives the patterns.
Academic/Practical Relevance: Many industries are increasingly using sensors to drive improvements in key performance metrics (e.g., asset uptime) through better information on operating conditions. Sensors will communicate amongst themselves to improve estimation. This IoT vision will have a major impact on operations management (OM), and OM scholars need to develop and examine models and frameworks to better understand sensor interactions.
Methodology: Analytic modeling combining decision-making, estimation, optimization, and learning.
Results: We show that when selecting its target(s) each sensor needs to consider both the measurement quality of the other sensors and its level of familiarity with their inference models, i.e., the precision of its beliefs. We establish that the state of the environment plays a key role in mediating quality and familiarity. When sensor qualities are public, we show that each sensor eventually settles on a constant target set but this long run target set is sample-path dependent (i.e., dependent on past states) and varies by sensor. The long run network, however, can be fully defined at time zero as a random directed graph and, hence, one can probabilistically predict it. This prediction can be made perfect (i.e., the network can be identified in a deterministic way) after observing the state values for a limited number of periods. When sensor qualities are private, our results reveal that sensors may not settle on a constant target set but the subset it cycles amongst can still be stochastically predicted.
Managerial Implications: Our work allows managers to predict (and influence) the set of other firms with which their sensors will form information links. Analogous to a manufacturer mapping its supplier base to help manage supply continuity, our work enables a firm to map its "sensor-based-information" suppliers to help manage information continuity.

Key words: Estimation; Information Sharing; Sensor Collaboration; Robust Optimization
History: Version: October 16, 2020

---

1. Introduction

The Internet of Things (IoT) and its sensor-based monitoring offers firms across a wide range of industries (including energy, healthcare, manufacturing, retail, and transportation) a vast opportunity to improve their operations through better information. Take the process industries, from upstream exploration through downstream commodity production, as an example. Asset uptime and worker safety are crucial operational metrics: a single day lost to unplanned downtime can cost a natural gas platform up to $25 million (Winnig 2016), and "on a global scale, unplanned shutdowns in the process industry cost 5 percent of total annual production—that's as much as

$30 billion a year" (GEPower 2018, p.2). Downtime is often caused by abnormal operating conditions; excessive vibration, for example, can be an important source or a leading indicator of equipment failure (FlukeCorp 2020), and an excessive ambient chemical concentration can cause a fire/explosion or be detrimental to worker health. Thus, process industries are increasingly turning to the IoT to improve asset uptime and promote worker safety by deploying sensors on equipment and workers to monitor key variables, e.g., vibration and trace chemicals, so as to help prevent equipment failure and other unsafe conditions (BehrTech 2020).[1]

An oil refinery depends on the correct functioning of tens of thousands of machines (Evans and Annunziata 2012), and "a typical oil drilling platform today might use 30,000 sensors, watching over the performance of dozens of systems" (McKinsey 2015, p.9). In their report, McKinsey claim that there is a vast untapped potential to use sensor information to make better operations decisions related to predictive maintenance and worker safety, and they comment that "the real value of IoT applications [comes] from analyzing data from multiple sensors" (p.104). Critical operating-condition variables such as vibration and trace chemical levels are challenging to measure and sensor estimates are inherently noisy. An important motive for sensor communication in an industrial worksite is to improve estimation of the operational variables of interest by sharing estimates across multiple sensors; that is, sensor communication can improve estimation. In turn, better estimation of key operating-condition variables enables more informed decisions and actions, e.g., when to dispatch technicians or remotely switch off/slow down certain machines if vibration levels are in a warning state or when to evacuate workers from a certain area if dangerous gases reach a certain level. These actions directly contribute to asset uptime and worker safety.

To appreciate the challenges and opportunities presented by communication between vibration sensors on different pieces of equipment or between chemical sensors on equipment and workers, for example, one must account for the scale and decentralized nature of process industry worksites. From a scale perspective, a large oil refinery complex can have more than a thousand workers and, as noted above, many thousands of critical equipment assets. While it would be theoretically desirable to have all vibration sensors (or chemical sensors) in the same local environment communicating with each other at each instant in time, it is not practically possible. Bandwidth, channel capacity, and/or energy consumption concerns place important limits on communication in sensor networks such that a sensor can only communicate with a subset of the other sensors in its environment at any given moment (Shi and Zhang 2012, Han et al. 2017, Wu et al. 2020). Therefore, a sensor

[1] Vibration is a key variable and "vibration analysis is the most commonly used condition monitoring technique" (Syed and Pai 2016).

cannot simultaneously solicit estimates from the tens, hundreds, or maybe even thousands, of other relevant sensors in its environment: it must choose at each moment a relatively small subset of the others sensors to target. This limitation might suggest that a sensor should target the most precise (highest quality) sensors in its environment, that is, those with the lowest measurement noise.

This intuition breaks down when one considers that "the oil and gas industry's business model and workforce include many stakeholders and partners. A single oil platform or refinery represents multiple companies; oil company employees work next to contractors, services organizations, and consultants. A range of manufacturers provides heavy equipment" (Winnig 2016, p.16) that often "sits side by side with competitors' equipment" (Winnig 2016, p.8). While estimation performance would be highest if firms shared all relevant information, sensor technology (e.g., inference model, measurement technique, raw readings) is often proprietary and closely guarded, and so companies are reluctant to share all sensor-related data. However, according to a GE Oil & Gas executive discussing sensor communication in operations, firms will engage in partial sharing of sensor-generated information as "they realize they can learn from each other" (Jernigan et al. 2016, p.6). For example, a vibration sensor from one company might share its current estimate with a sensor from another company but not share its actual reading or its underlying inference algorithm. Therefore, when choosing at each moment which other sensors to target to improve its own estimate, a sensor must account not just for the quality of other sensors (which it may or may not know) but also for its limited understanding of how the other sensors form their estimates.

There are at least two important aspects to the process-industry vision of using widely deployed sensors to improve operational performance: (a) generating better operating-condition information through sensor communication, and (b) enabling smarter decisions and actions through this better information. These two aspects are somewhat related but also quite decoupled because better decisions and actions are enabled by better estimation irrespective of the specific decision-action context. In this paper we focus on (a) — the sensor communication aspect — and we seek to answer the following question: in a dynamic environment populated by sensors from different firms, each trying to maximize its own estimation performance, which subset of other-firm sensors should a sensor target at each point in time (given a limit on the number targeted); and how does that subset evolve over time (when one accounts for the fact that receiving estimates from other sensors can enable a sensor to learn something about the inference models used by these sensors)? We discuss (b) as a direction for future research in the conclusion of this paper.

We have adopted the process industry as a motivating context but there are many other settings in which noisy sensors from multiple firms operate in a common environment and can therefore benefit from communication. Indeed, the concept of information "fusion across sensors [that] nominally measure the same property [to] reduce or eliminate noise and errors" (Mitchell 2007, p.4-5) is envisioned as a remedy for sensor quality concerns in a wide range of IoT settings including energy, environmental monitoring, infrastructure, and various other industrial applications, and systems of sensors "will not necessarily be owned by one group [and] more often than not, more than one organization will operate in the same system" (IEC 2015, p.33). This is true even in healthcare settings; for example, many firms are developing minimally-invasive wearable bio-sensors to measure blood glucose levels and it is envisioned that a diabetic person might simultaneously wear multiple sensors (Xiong et al. 2011). More broadly, the emerging concept of Industry 4.0 — the marriage of the digital and physical worlds in manufacturing and supply chain systems (Lennon Olsen and Tomlin 2019) — will depend on sensors that "connect and communicate with one another, ... and decentralized decision making — the ability of cyber-physical systems to make simple decisions on their own and become as autonomous as possible" (Marr 2016).

With autonomous IoT sensors communicating across companies, firms will "become enmeshed in a network of organizational relationships that require dedicated resources and management attention" (Jernigan et al. 2016, p.14). Therefore, analogous to the current interest that consumer, media, and technology firms have in understanding social network formation (Momot et al. 2019), industrial firms will increasingly seek to understand the formation and evolution of autonomous inter-firm IoT sensor communication patterns as Industry 4.0 matures. The goal of this paper is to take the first steps in developing such an understanding. If one considers humans to be fully rational (as sensors are), then our work can be applied to a wisdom-of-the-crowd, social-network problem in which individuals forecasting a dynamically evolving variable (e.g., product demand) choose over time which other individuals from whom to solicit forecasts when individuals do not know but can learn about the forecast models used by other individuals. Full rationality is arguably a strong assumption for human forecasters (Tong and Feiler 2017). Therefore, our paper applies more directly to information fusion in decentralized, i.e., multi-firm, sensor networks.

In this paper, we consider a collection of autonomous sensors (i.e., no central governing entity) operating in a common environment in which a state variable (representing an operating condition such as vibration) evolves according to an autoregressive time series model. Each sensor is unbiased but imperfect and generates a private, zero-mean noisy signal of the state in each time period.

After observing its own private signal of the state in a period, each sensor chooses a subset of other sensors to target (i.e., from which sensors to solicit state estimates) so as to generate an improved estimate. Sensors do not know the inference models used by other sensors for estimation but can learn about them over time if they communicate. In our base model, we assume that sensors know the qualities of all other sensors. We then relax the known-quality assumption and adopt a robust optimization approach in which the knowledge about other sensor qualities is ambiguous. For expositional clarity, we focus on the setting in which a sensor can choose only one target during each period but in the appendix show how all our results naturally extend to the setting in which multiple simultaneous targets are allowed.

Among other results, we establish that the updating of the sensor's state estimate depends on (a) the qualities of both the targeted and the targeting sensors; and (b) the targeting sensors beliefs about the targeted-sensor's inference model, which update over time the more that sensor is targeted, that is, the more "familiar" a sensor becomes with another sensor's inference model parameters. We show that the state of the environment plays a key role in determining the weights placed on quality and familiarity when selecting a target in any given period. We prove that when qualities are known and asymmetric, each sensor will eventually target a single sensor in all future periods but this long run contact can vary by sensor. State dependency means that these long run contacts are sample-path dependent, and hence, even for each particular sensor, the long run contact can vary depending on the realization of state over time. Nevertheless, we demonstrate that the long run communication network that forms between sensors can be fully defined at time zero as a random directed graph, which means that one can probabilistically predict the long-run communication patterns that will emerge. Furthermore, this prediction can be made perfect (i.e., deterministic as opposed to probabilistic) after observing the state values for a limited number of periods. When qualities are not common knowledge, i.e., sensors face ambiguity with respect to other sensor qualities, randomizing across some subset of sensors may be optimal in the long run, but we provide an intuitive sufficient condition under which a deterministic targeting policy is optimal.

These technical results have a number of important managerial implications. In particular, we note that sensor communication enables improved estimation but it creates a dependency on other firms' sensors. Consider a multi-firm worksite in the process industry for example. A firm's asset uptime will depend to some degree on sensors from other firms if that firm's sensor targets those other sensors to improve its own vibration estimate. Therefore, those other firms are "information suppliers" to the firm. Our work help a firm identify which other firms have a high likelihood of being its long-run information suppliers. Just as a manufacturer maps its supply base and monitors events (e.g., bankruptcies) that can impact supply continuity, a firm can map its important sensor-based-information suppliers and monitor for events that might affect information continuity, e.g., a corporate ownership change of an important group of targeted sensors that might negatively affect information sharing. Our results also establish the key levers that a firm can use to influence which other firms will be its long-run contacts (information suppliers). Knowing these levers, a firm can then exert influence over the set of firms on which it depends.

The rest of the paper is organized as follows. The most relevant literature is discussed in §2. The base model is described in §3. Analysis and results are presented in §4 and §5. The extension to unknown sensor quality is developed and analyzed in §6. A number of other extensions are summarized in §7 and fully developed in the appendices. Conclusions are discussed in §8. Appendices A through D are available in the e-companion. Appendices E through I are available from the authors upon request.

2. Related Studies

Our research is related to a number of streams of literature that examine information sharing for the purpose of improved estimation or forecasting.

The idea of sharing information between sensors is not a new one. Multi-sensor data fusion, defined by Mitchell (2007)[p.3.] as "the theory, techniques and tools which are used for combining sensor data ... so that it is, in some sense, better than would be possible if the data sources were used individually" emerged as a problem domain in the 1990s due to the U.S. military's desire to enable more-complete or higher-quality surveillance of geographic areas. It has since grown to encompass diverse applications in artificial intelligence, robotics, and environmental, equipment and health monitoring. Sensor fusion is typically focused on developing efficient and effective data architectures, processing techniques and protocols for aggregating information collected by a defined network of sensors (Hall and Llinas 1997, Mitchell 2007, Khaleghi et al. 2013). Recognizing that bandwidth, channel capacity, and/or energy consumption concerns place important limits on sensor communication, there is a stream of sensor-fusion work related to sensor scheduling in which a sensor can only communicate with a limited number of other sensors in a given period (Shi and Zhang 2012, Vitus et al. 2012, Yang et al. 2014, Han et al. 2017, Wu et al. 2020).

The sensor fusion and sensor scheduling literatures presuppose a collection of sensors deployed by a single governing entity such that all sensors are willing to share all relevant information; the challenge is to efficiently communicate and aggregate the information. Our work differs significantly in intent from that literature and is distinguished by its focus on settings in which autonomous sensors are deployed by different firms without a central governing entity. In such settings, sensors do not have full information on the estimation approaches of other sensors but can learn about them over time.

Forecasting is a central concern in operations management, and it has long been recognized that combining demand estimates/information from multiple individuals or firms can improve forecast accuracy (e.g., Fisher and Raman 1996, Swaminathan and Tayur 2003, Gaur et al. 2007, Simchi-Levi 2010). More recently, motivated by the emergence of external and internal prediction markets, Bassamboo et al. (2018) empirically explores the effect of group size on forecast accuracy, finding that aggregation across larger groups improves accuracy. The notion that aggregation of a large number of estimates can improve estimation—sometimes described as the wisdom of crowds—has also received significant attention in the decision analysis, economics, forecasting, social network, and other literatures (e.g., Bates and Granger 1969, Ashton and Ashton 1985, Winkler and Clemen 2004, Wallis 2011, Acemoglu et al. 2014a,b, Atanasov et al. 2016, Tsoukalas and Falk 2019). Through that lens, one can view our work as exploring a related but different question: when each individual in the crowd wants to improve his or her own estimate (but cannot ask everyone in the crowd) then who in the crowd should an individual target?

With that lens in mind, the paper most related to our work appears to be Sethi and Yildiz (2016) who examine communications between human experts that independently observe a static white-noise process. In each period, each expert estimates the current state with some randomly-drawn precision (i.e., quality) whose realization is publicly observable to all experts. These human experts may differ in their private opinions on the mean level of the process. Each expert can solicit an estimate from one other expert in each period. The authors examine the types of long run communication networks that can emerge. Although sharing certain features (e.g., target selection must tradeoff between quality and unknown beliefs), our work differs significantly from Sethi and Yildiz (2016) in some fundamental aspects that are driven by our IoT-sensor motivating context. For example, we consider a dynamic (not static) environment because that is a typical feature of the environments in which sensors are deployed. We establish the importance of this distinction by proving that–different to a static random environment–the state and its dynamics are a crucial driver of target selection. Furthermore, the human experts' qualities are randomly redrawn in every period in Sethi and Yildiz (2016), with realizations being common knowledge. This highlights two other critical differences in our work driven by the IoT context: sensor qualities are not typically random and, more importantly, sensor qualities may not be known to other sensors. To accommodate this unknown-quality reality, we adopt a robust optimization framework in which sensor qualities are ambiguous and target selection needs to be robust to this ambiguity.

Finally, our work is also related to the general theory of robust optimization and estimation; relevant papers from the operations literature include Liyanage and Shanthikumar (2005), Perakis and Roels (2008), Delage and Iancu (2015), Saghafian and Tomlin (2016), Mišić and Perakis (2019), and references therein. For some general theoretical results on the percentile optimization approach that we utilize, we refer interested readers to Nemirovski and Shapiro (2006), Delage and Mannor (2010), Bren and Saghafian (2019), and references therein.

3. The Model

We first present a high level description before formalizing the model. We consider a collection of sensors deployed by different firms in a common environment. Each sensor estimates (using its own measurement and inference model) a state variable representing the operating condition in the environment, e.g., the vibration level. The objective of each sensor is to generate the most accurate state estimate it can. We intentionally omit any resulting operational action because there is a wide class of action problems for which the goal of the sensor is to generate the best state estimate it can. Our model is in fact indifferent to the action problem so long as action decisions benefit from higher-quality state information. This means that the model we study is quite general.[2]

Sensors generate quick but noisy estimates in their local processing units. In practice, these noisy estimates are sometimes augmented by more accurate but less frequent or slower estimates. For example, maintenance technicians might periodically record “vibration data with a handheld data collector” (IKM 2019) that can then be fed back to the sensor.[3] Alteratively, in addition to its own rapid, local (“edge”) computation, a sensor might also “offload” its data to a remote computer or the cloud for more computationally-intensive and more accurate estimation but this comes at the expense of a communications delay (Ran et al. 2017, Ballotta et al. 2019, Xu et al. 2020). In our base model, we assume that each sensor has a slower-but-more-accurate estimation approach in place (e.g., technician inspection or remote offloading) and for simplicity assume that it incurs

[2] We make no assumption that the devices associated with the sensors are even engaged in related or analogous actions. We merely assume that each sensor’s objective is to generate the highest quality state estimate it can for its associated device.

[3] Analogously in the diabetic monitoring context, it is thought that wearable blood-glucose (BG) biosensors will require “frequent calibration against direct BG data” obtained by precise but invasive means (Chen et al. 2017, p.8).

a delay of one period and is perfectly accurate. In Appendix C, we analyze relaxations in which the slower-but-more-accurate approach has a general delay and may or may not provide perfect estimation and we also consider relaxations in which the slower-but-more-accurate approach does not exist. We show that our key results extend to these relaxations.

We consider a partial-information sharing regime in which the sensor-owning entities are willing to share some but not all information. In particular, each sensor is willing to share its current state estimate and possibly its underlying sensor quality but not its inference model or raw measurement. A sensor can solicit estimates from other sensors to improve its own estimate but it is limited in the number of other sensors it can target because of the communication constraints discussed earlier in the paper. We explore the problem of determining for each sensor (in each period) which other sensor(s) it should target so as to most improve the accuracy of its own state estimate.

In what follows, we formally describe the environment, individual sensor measurement and state estimation, sensor collaboration, and finally the target selection problem whereby each sensor chooses from which other sensors to solicit state estimates.

Environment: A collection $\mathcal{N} \triangleq \{1, 2, \cdots, n\}$ of autonomous sensors exist in a common environment that is defined by an operating-condition variable $S \in \mathbb{R}$ (e.g., vibration) whose discrete-time state evolution is governed by a first-order autoregressive (AR(1)) process:

$$S_t = \alpha + \beta S_{t-1} + \tilde{\epsilon}_t \tag{1}$$

for $t = 1, 2, \cdots, \infty$, where $\tilde{\epsilon}_t$ are i.i.d. normal white-noise random variables with mean 0 and variance normalized to 1. We note that autoregressive behavior is a common phenomenon and that AR models are used to estimate a wide range of dynamic properties including two of our motivating examples: equipment vibration (Thanagasundram and Schlindwein 2006, Ayaz 2014) and blood glucose (Sparacino et al. 2007, Leal et al. 2010). We adopt an AR(1) model for reasons of parsimony and note that such a model has been adopted in the sensor network literature, e.g., Vitus et al. (2012), Shi and Zhang (2012) and Ballotta et al. (2019).[4]

Individual Sensor Measurement and State Estimation: At the beginning of each time period t, each sensor $i \in \mathcal{N}$ privately generates a noisy signal (observation) Γ_{it} of the state variable S_t. In many IoT settings, e.g., when the variable-of-interest is difficult or time-consuming to measure, this signal is indirectly generated by measuring some other related properties and mapping these measurements into the variable-of-interest. Different sensor technologies may rely on different indirect

[4] We acknowledge that higher-order models can have more predictive power; however, it is a generally advised principle that one should select a model of minimum order needed for a good fit, and AR(1) models are sometimes used for estimation, e.g., Sparacino et al. (2007).

properties, and hence, different mappings. To avoid unnecessary notational burden, we suppress the raw readings and related mapping, and instead focus on the final noisy signal of the current state (S_t) privately derived by sensor i:

$$\Gamma_{it} = S_t + \epsilon_{it}, \tag{2}$$

where ϵ_{it} are i.i.d.[5] normal white noises with mean 0 and variance $1/(q_i)^2$, with q_i representing the quality of sensor i. That is, a higher quality sensor has a higher precision.

Each sensor $i \in \mathcal{N}$ knows that the environment evolves according to an AR(1) process but does not know the true parameters of the AR(1) process. Specifically, when using its signal to estimate the current state of the environment, sensor i uses its own inference model—developed based on its firms's training algorithms and data sets prior to deployment—which is given by

$$S_{it} = \hat{\alpha}_i + \hat{\beta}_i S_{t-1} + \tilde{\epsilon}_t, \tag{3}$$

where $\hat{\alpha}_i$ and $\hat{\beta}_i$ are sensor i's estimates of the process parameters α and β. Immediately before period t starts, the realized value of the previous period state s_{t-1} is revealed to each sensor (because of our base-model assumption that the slow-but-more-accurate estimation approach of each sensor is perfect with a one-period delay), and the system moves to the next period.[6]

At the beginning of each period t, knowing the realization of the previous period state s_{t-1}, but prior to receiving the noisy signal Γ_{it}, sensor i believes (based on its inference model (3)) that the current state S_t follows a normal distribution with mean $\hat{\alpha}_i + \hat{\beta}_i s_{t-1}$ and variance 1. Upon realizing the current signal $\Gamma_{it} = \gamma_{it}$, sensor i updates its prior belief about the current state according to Bayes' rule. Since both the signal received about the state and the prior on state have a normal distribution (see (2) and (3)), it follows from Bayes' rule that sensor i's posterior belief about the state is also normally distributed but with a mean and variance given by

$$\mathbb{E}[S_{it} | \Gamma_{it} = \gamma_{it}] = \frac{\hat{\alpha}_i + \hat{\beta}_i s_{t-1}}{1 + q_i^2} + \frac{q_i^2}{1 + q_i^2} \gamma_{it} \tag{4}$$

and

$$Var[S_{it} | \Gamma_{it} = \gamma_{it}] = \frac{1}{1 + q_i^2}, \tag{5}$$

[5] In Appendix D, we study scenarios in which the measurement errors are correlated among sensors, and show that our main results extend to such scenarios. However, even when these measurement errors are not correlated, it should be noted that sensors' readings given in (2) are still correlated both within and across periods.

[6] Our base-model analysis and findings immediately extends to a setting in which state realization occurs less frequently (e.g., every $T > 1$ periods because technician inspections or data offloading is periodic due to travel or communication burdens) but target selection remains constant between realizations; this only requires a re-scaling of time, i.e., changing the definition of a period. Also, please see Appendix G for generalizations in which the slow-but-more-accurate estimation is not perfect and/or has a general delay of more than one period.

respectively. The higher the quality of sensor i, the more weight it places on its signal when updating its mean belief, and the larger the associated variance reduction.

Information Sharing and Sensor Collaboration: Each sensor $i \in \mathcal{N}$ is aware of all the other sensors in the environment. All sensors in the collection $\mathcal{N}$ are willing to collaborate in the following manner: in each period t, after all sensors have formed updated beliefs based on their private signals (according to (4) and (5) above), any sensor $j \in \mathcal{N}$ is willing to share its best estimate of state (according to the expected value of the squared error loss) which is its updated mean prediction of state $E[S_{jt}|\Gamma_{jt} = \gamma_{jt}]$ with any other sensor i that requests it.[7] Sensors are deployed by different firms, and therefore sensor $i \in \mathcal{N} \setminus \{j\}$ may not know the inference model parameters $\hat{\alpha}_j$ and $\hat{\beta}_j$ used by sensor j because firms typically train their sensors (pre-deployment) differently using different algorithms and training data sets that are often privately owned. We assume that at time $t = 0$ sensor i believes that sensor j's inference model parameters $\hat{\alpha}_j$ and $\hat{\beta}_j$ come from independent normal distributions $N(\hat{\alpha}_j, 1/v_{ij0}^2)$ and $N(\hat{\beta}_j, 1/w_{ij0}^2)$, respectively.[8] In this setting, parameters $v_{ij0} > 0$ and $w_{ij0} > 0$ represent sensor i's initial *familiarity* with sensor j's inference model. Higher familiarity values indicate more precise beliefs. A setting in which sensor i fully knows sensor j's inference model parameters can be obtained by setting $v_{ij0} = w_{ij0} = \infty$. To gain insights, in our base model we assume that sensor qualities q_i are common knowledge to all $i \in \mathcal{N}$, but this is relaxed in §6.

Target Selection: In each period t, after updating its state estimate based on its private signal as in (4) and (5) above, each sensor i chooses a set of sensors from which to request state estimates, i.e., their updated mean beliefs about the state. We do not model the actions of devices associated with sensors but implicitly assume that the action payoff is increasing in the quality of the state estimate. Thus, in choosing which sensors to target, sensor i selects those sensors that will most improve its own estimate. By most improvement, we mean that sensor i's resulting updated state distribution gives the lowest expected squared error of estimation.[9] In particular, sensor i's decision in each period t is based on the following optimization problem:

$$\min_{\tilde{s}_{it} \in \mathbb{R}, \mathbf{a}_{it} \in \{0,1\}^{n-1}} \quad \mathbb{E}_{S_t \sim F_{it}^{\mathbf{a}_{it}}} \left[\tilde{s}_{it} - S_t \right]^2 \tag{6}$$

[7] Our work easily extends to a setting in which j will only collaborate with some subset of $\mathcal{N}$.

[8] The extension to a setting in which the means of these normal distributions are not correct is contained in Appendix B which establishes that the key results in the paper still hold.

[9] As we will show, this implies that each sensor targets those sensors that provide it with the most information about the state. We use the expected value of the squared error as our targeting objective function mainly because it is a common loss function used in the literature of Machine Learning and Estimation Theory. However, all our results hold for any targeting objective function that is strictly increasing in the expected squared error of estimation.

$$s.t.$$

$$0 < c \sum_{j \in \mathcal{N} \setminus \{i\}} a_{ijt} \leq b,$$

where the vector $\mathbf{a}_{it} \in \{0,1\}^{n-1}$ is composed of elements a_{ijt} with $a_{ijt} = 1$ if i targets j at time t, and $a_{ijt} = 0$ otherwise, $F_{it}^{\mathbf{a}_{it}}$ is the posterior distribution of sensor i's belief about the state after communicating with the selected targets at time t, c is the cost of communication per target in each period, and b is a communication budget in each period. Thus, in (6), sensor i decides upon the vector $\mathbf{a}_{it} \in \{0,1\}^{n-1}$ (who to communicate to) subject to its communication budget. This decision impacts the sensor's posterior belief about the state (once the communication is made), and hence, the sensors is trying to communicate to sensors that will allow for the most desired posterior belief about the state, where most desired is measured in terms of ℓ_2-norm loss in estimation.

We define $k = b/c$, and refer to it as the targeting channel capacity because $\lfloor k \rfloor$ represents the maximum number of targets from which a sensor can solicit estimates in a period. For expositional ease, we focus on the case where a sensor can choose only one target in each period, i.e., $\lfloor k \rfloor = 1$, because this is when choosing the right target is most important. Our results readily extend to the case of $\lfloor k \rfloor > 1$ as we mention at times in the paper and fully show in Appendix A.

Other Extensions: In addition to the relaxations mentioned during the model description (i.e., delayed and imperfect state realizations, correlation in sensors' errors, and incorrect beliefs about the means of other sensors' inference parameters), we extend our base model to consider non-myopic sensors that care about estimation in future periods and also extend it to consider a setting in which sensors update their own inference models over time in a Bayesian fashion. All these extensions and their results are summarized in §7 but we relegate their detailed analysis and discussion to the appendices for brevity.

4. Preliminaries: Targeting Equivalence and Familiarity Dynamics

As a preliminary to our exploration of how sensor communications evolve over time, we first develop an equivalent target-selection problem and analyze how any given sensor's beliefs about other sensors' parameters update from one period to the next.

Targeting Equivalence: We begin by establishing that the target selection problem in (6) above is equivalent to one in which sensor i selects as its target the sensor that provides i with the most informative signal about the current state, where a less noisy signal (i.e., one with a lower variance) is more informative.[10] Importantly, we will show that the informativeness of a signal

[10] Note that the information entropy of any normally distributed random variable depends only on its variance.

depends not only on the quality of the potential target sensor j, but also on the receiving sensor i's familiarity with sensor j's inference model. In particular, given its privately generated signal Γ_{jt} in period t, sensor j provides sensor i with its best current estimate of state, which is $\mathbb{E}[S_{jt}|\Gamma_{jt}]$, i.e., its updated/latest expected belief about the current state S_t. Now, from sensor i's perspective, $\mathbb{E}[S_{jt}|\Gamma_{jt}]$ is formed according to:

$$\mathbb{E}[S_{ijt}|\Gamma_{jt}] = \frac{\hat{\alpha}_{ijt} + \hat{\beta}_{ijt}s_{t-1}}{1+q_j^2} + \frac{q_j^2}{1+q_j^2}\Gamma_{jt}, \tag{7}$$

which is similar to (4) but where $\hat{\alpha}_{ijt}$ and $\hat{\beta}_{ijt}$ reflect sensor i's beliefs at time t about sensor j's inference parameters $\hat{\alpha}_j$ and $\hat{\beta}_j$. (The relevant dynamic updating mechanism is developed below.) Because $\Gamma_{jt} = S_t + \epsilon_{jt}$ from (2), this value $\mathbb{E}[S_{ijt}|\Gamma_{jt}]$ provides sensor i with the following noisy signal regarding the state S_t:

$$\frac{1+q_j^2}{q_j^2}\mathbb{E}[S_{ijt}|\Gamma_{jt}] = S_t + \epsilon_{jt} + \frac{\hat{\alpha}_{ijt} + \hat{\beta}_{ijt}s_{t-1}}{q_j^2}. \tag{8}$$

We denote the variance in this signal's noise as:

$$\sigma_t^2(i,j,s_{t-1}) = Var\big[\epsilon_{jt} + \frac{\hat{\alpha}_{ijt} + \hat{\beta}_{ijt}s_{t-1}}{q_j^2}\big]. \tag{9}$$

There are two independent sources of noise in this signal: (a) the inherent white noise ϵ_{jt} in sensor j's measurement Γ_{jt} (which has a variance of $1/q_j^2$), and (b) the noise caused by sensor i's lack of familiarity with sensor j's inference model. For notational convenience, we define the random variable $\Xi_{ijt}(s_{t-1}) = \hat{\alpha}_{ijt} + \hat{\beta}_{ijt}s_{t-1}$, where its dependence on the prior state value s_{t-1} is explicitly noted. Defining the precision $\psi_{ijt}(s_{t-1}) \triangleq 1/Var\big[\Xi_{ijt}(s_{t-1})\big]$, it follows from (9) that

$$\sigma_t^2(i,j,s_{t-1}) \;=\; \frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4}. \tag{10}$$

Under a variance reduction objective, sensor i's target in period t is

$$\begin{aligned} j_{it}^* &\triangleq \arg\min_{j\in\mathcal{N}\setminus\{i\}} \sigma_t^2(i,j,s_{t-1}) \\ &= \arg\min_{j\in\mathcal{N}\setminus\{i\}} \Big\{\frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4}\Big\}. \end{aligned} \tag{11}$$

The following result establishes that the original target selection problem in (6) is equivalent to the variance reduction target selection (11); that is, both objectives result in the same target.[11]

PROPOSITION 1 **(Target Selection and Variance Reduction)**. *If channel capacity $\lfloor k \rfloor = 1$, then under (6), $a_{ijt}^* = \mathbb{1}_{\{j=j_{it}^*\}}$, where j_{it}^* is given by (11), and $\mathbb{1}_{\{\cdot\}}$ is the indicator function.*

[11] Without loss of generality, we assume ties in (6) and (11) are broken by choosing the sensor with the lower index.

This result readily extends to a general channel capacity $k \geq 2$: each sensor i will select the $\lfloor k \rfloor$ other sensors that provide the lowest variance of signal. That is, it chooses the $\lfloor k \rfloor$ most-informative sensors (from its perspective) and solicits their state estimates; see Appendix A for more details and extension of our results with $k \geq 2$. This *rank-ordering structure* also illuminates the importance of studying the case where each sensor can only target one sensor ($\lfloor k \rfloor = 1$) during each period (while targets may vary across periods). As noted earlier, this is because $\lfloor k \rfloor = 1$ represents the scenario in which choosing the right target is most critical.

Recall that, by definition, the variance of the random variable $\Xi_{ijt}(s_{t-1}) = \hat{\alpha}_{ijt} + \hat{\beta}_{ijt} s_{t-1}$ is given by $1/\psi_{ijt}(s_{t-1})$. In what follows, we therefore refer to $\psi_{ijt}(s)$ as the *familiarity function* that sensor i has for sensor j at time t, and we refer to $\psi_{ijt}(s_{t-1})$ as the *familiarity value*, i.e., the familiarity function evaluated at the latest realized state $s = s_{t-1}$. We use the term familiarity to convey the notion that higher values imply less noise in sensor i's beliefs about sensor j's underlying inference model. Importantly, as we will establish below, sensor i does not need to separately update its beliefs over time about parameters $\hat{\alpha}_{ijt}$ and $\hat{\beta}_{ijt}$ of sensor j's inference model; it suffices to update the familiarity function $\psi_{ijt}(s)$.

Familiarity Dynamics: To operationalize the target selection problem (11), we now examine how any given sensor's familiarity function with respect to some other sensor evolves over time. In particular, we develop the mechanism through which the time-t familiarity function is updated to that at time $t+1$, i.e., how $\psi_{ijt}(s)$ updates to $\psi_{ij,t+1}(s)$. To that end, we first note that it follows from the definition $\psi_{ijt}(s_{t-1}) \triangleq 1/Var\big[\Xi_{ijt}(s_{t-1})\big]$ that the initial familiarity function is given by

$$\psi_{ij1}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2}. \tag{12}$$

We also note that, if sensor i communicates with sensor j at time t, then it follows from (7) that i receives the following signal about the random variable $\Xi_{ijt}(s_{t-1}) = \hat{\alpha}_{ijt} + \hat{\beta}_{ijt} s_{t-1}$:

$$(1+q_j^2)\mathbb{E}[S_{ijt}|\Gamma_{jt}] = \Xi_{ijt}(s_{t-1}) + q_j^2(S_t + \epsilon_{jt}). \tag{13}$$

There are two independent sources of noise in this signal: (a) the noise in sensor i's own estimate of the current state S_t, which has variance of $1/(1+q_i^2)$ (see (5)), and (b) the inherent white noise ϵ_{jt} in sensor j measurement, which has a variance of $1/q_j^2$. Thus, based on (13), the variance in the signal's noise is given by $Var[q_j^2(S_t + \epsilon_{jt})] = q_j^4/(1+q_i^2) + q_j^2$. Using Bayesian updating of the aggregated univariate random variable $\Xi_{ijt}(s_{t-1})$ after observing this signal, which eliminates the need to explicitly update and carry over the joint distribution of $\hat{\alpha}_{ijt}$ and $\hat{\beta}_{ijt}$ (hence, their covariance matrix), we can show the following result.[12]

PROPOSITION 2 **(Familiarity Dynamics)**. *For any $s \in \mathbb{R}$:*

(i) $\psi_{ij,t+1}(s) = \psi_{ijt}(s) + \delta(q_i, q_j, a_{ijt})$*, where* $\delta(q_i, q_j, a_{ijt}) \triangleq f(q_i, q_j) a_{ijt}$ *and*

$$f(q_i, q_j) \triangleq \frac{(1+q_i^2)}{q_j^2(1+q_i^2+q_j^2)}. \tag{14}$$

(ii) For all $t = 1, 2, 3, \cdots$*, we have*

$$\psi_{ij,t+1}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j) \sum_{l=1}^{t} a_{ijl}. \tag{15}$$

Intuitively, sensor i's familiarity function for sensor j changes from time t to time $t+1$ if, and only if, i targets j at time t, i.e., $a_{ijt} = 1$. Moreover, if i targets j then the gain in i's familiarity with j does not depend on the state: the gain is given by $f(q_i, q_j)$, which we refer to as the *stickiness factor*. It is noteworthy, however, that the gain depends on both the sender's (j's) and the receiver's (i's) qualities. More importantly, it follows from (15) that to calculate the current familiarity function that sensor i has for sensor j we only need to know the number of times that i selected j as its target; we do not need to know in which periods those selections occurred. This is primarily due to the fact that when i communicates to j in any given period the gain in i's familiarity with j depends only on i's ability to interpret the signal from j about the state. This ability to interpret depends only on the time-invariant qualities of the receiver q_i and the sender q_j and does not depend on the state which is time varying.

5. Communication Networks: Who Targets Who?

With the equivalent target selection problem and familiarity dynamics developed, we now characterize how target selection evolves over time. In choosing a target in period t, any given sensor i needs to consider both the quality of each other sensor j and its own current familiarity value $\psi_{ijt}(s_{t-1})$ for each sensor j; see the targeting criterion (11). That is, beside quality, the attractiveness of j as a potential target for i depends on the familiarity value $\psi_{ijt}(s_{t-1})$. The familiarity value, in turn, depends explicitly on the previous state s_{t-1} but also implicity on all prior states through their influence on prior targeting of sensor j by sensor i. Thus, target selection in each period depends on the history of state realizations up to that period.[13]

[12] We also provide a second proof of this result in which we explicitly use the joint distribution of $\hat{\alpha}_{ijt}$ and $\hat{\beta}_{ijt}$ and characterize how their covariance matrix is updated over time; see proofs in Appendix G.

[13] This state dependency does not arise if the underlying environment is governed by a static i.i.d. white noise process, i.e., when $\beta = 0$. In that case, it follows directly from the above analysis that the familiarity function $\psi_{ij,t+1}(s) =$

5.1. Initial Target Selection

It is informative to first consider target selection at time $t=1$, because this initial selection highlights a key ongoing tradeoff between sensor qualities and state value. Consider any given sensor i, and assume (without loss of generality) that it can select its target from two sensors: a high-quality sensor (labeled h) and a lower quality sensor (labeled l). When should sensor i target sensor h? When should it target sensor l? How does the choice depend on the initial state s_0?

To answer these questions, let $r \triangleq q_l/q_h$ denote the quality ratio of sensors l and h. By definition, $0<r\leq 1$. Using (11) and (12), it follows that sensor i strictly prefers targeting the lower quality sensor (l) if, and only if,

$$\frac{(r\,q_h)^2+(\frac{1}{v_{il0}})^2+(\frac{s_0}{w_{il0}})^2}{(r\,q_h)^4} < \frac{q_h^2+(\frac{1}{v_{ih0}})^2+(\frac{s_0}{w_{ih0}})^2}{q_h^4}, \tag{16}$$

where $1/v_{ij0}^2$ and $1/w_{ij0}^2$ are sensor i's initial belief variances about sensor $j\in\{h,l\}$ parameters $\hat{\alpha}_j$ and $\hat{\beta}_j$ respectively. From (16), it can be seen that i strictly prefers to target l if, and only if,

$$c_0+c_1 s_0^2>0, \tag{17}$$

where

$$c_0=(r\,q_h)^2(r^2-1)+(\frac{r^2}{v_{ih0}})^2-(\frac{1}{v_{il0}})^2, \tag{18}$$

and

$$c_1=(\frac{r^2}{w_{ih0}})^2-(\frac{1}{w_{il0}})^2. \tag{19}$$

We note that c_0 reflects a tension between the difference in sensor qualities and the differences in i's initial familiarity with the inference model parameters $\hat{\alpha}_h$ and $\hat{\alpha}_l$. Similarly, c_1 reflects a tension between the difference in sensor qualities and the differences in i's initial familiarity with the inference model parameters $\hat{\beta}_h$ and $\hat{\beta}_l$. As (17) shows, state influences initial target selection, and this is through the c_1 term. The following result presents the conditions under which sensor i strictly prefers to target the lower quality sensor. Thus, it also sheds light on conditions under which sensor i prefers to sacrifice quality for familiarity. By an appropriate swapping of labels l and h, it can also be used to highlight conditions under which sensor i strictly prefers to target the higher quality sensor.

$v_{ij0}^2+f(q_i,q_j)\sum_{l=1}^{t}a_{ijt}$. This is independent of the state s, and therefore, target selection is sample path independent. From this perspective, one can view Proposition 2 as generalizing the belief updating expressions (7)-(9) in Sethi and Yildiz (2016) to the case of an AR(1) process.

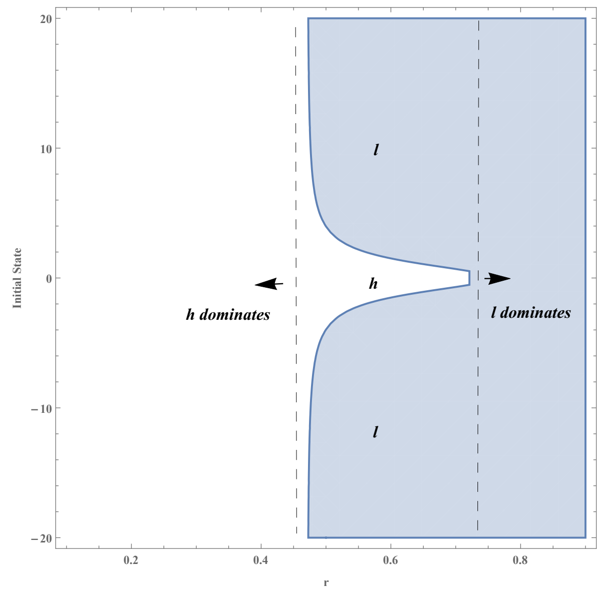

Figure 1 Initial selection between a higher quality sensor (h) and a lower quality sensor (l). The darker (lighter) area represents the region in which l (h) is the preferred sensor [Proposition 3 (i) applies in the intermediate region between the dashed lines.]

Proposition 3 **(Initial Selection)**. *A sensor i strictly prefers to target a lower quality sensor (l) than a higher quality one (h) at $t=1$ if, and only if, one of the following conditions holds:*
(i) $c_0 \leq 0$, $c_1 > 0$, and $|s_0| > \sqrt{-c_0/c_1}$,
(ii) $c_0 > 0$ and $c_1 < 0$, and $|s_0| \leq \sqrt{-c_0/c_1}$, or
(iii) $c_0 > 0$ and $c_1 \geq 0$.

This proposition highlights the interconnected roles that (a) sensor qualities, (b) familiarities, and (c) state play in target selection. Intuitively, if sensor i has more familiarity with the high-quality sensor's inference model parameters (i.e., $v_{ih0} \geq v_{il0}$ and $w_{ih0} \geq w_{il0}$), then the high-quality sensor is the inherently more attractive target regardless of state. This is reflected in the above proposition by the fact that $c_0 < 0$ and $c_1 < 0$ in this case and, therefore, sensor h is preferred. On the other hand, if sensor i has more familiarity with at least one of the low-quality sensor's inference parameters, then the high-quality sensor might not be the preferred target, because its estimate may prove to be more noisy from i's perspective (than the low-quality sensor). This tradeoff between quality and familiarity depends on the state (parts (i) and (ii) of Proposition 3) unless the familiarity advantage of the lower quality sensor compared to the higher quality one is so large that it makes the lower quality sensor the preferred target regardless of the state (part (iii) of Proposition 3).

As the quality ratio r increases from 0 to 1 (all else held constant) there are at most three distinct regions of target selection that identify the role of state, as illustrated in Figure 1.[14] When the quality ratio r is low, i.e., sensor h is of much higher quality than l, then h dominates l, i.e., h is targeted in all states. This h-dominating region always exists, but it does not cover the entire range $0 < r \leq 1$ unless $v_{ih0} \geq v_{il0}$ and $w_{ih0} \geq w_{il0}$, i.e., the high-quality sensor is the more familiar for both parameters. In contrast, when the quality ratio is high, i.e., sensor qualities relatively similar, then l dominates h, i.e., l is targeted in all states. This l-dominating region exists if, and only if, $v_{ih0} < v_{il0}$ and $w_{ih0} < w_{il0}$, i.e., the low quality sensor is the more familiar for both parameters. Importantly, there is an intermediate range of the quality ratio r (that extends to $r = 1$ if the familiarity ranking differs across v and w) in which state matters and the indifference curve $|s_0| = \sqrt{-c_0/c_1}$ completely characterizes target selection. Figure 1 illustrates an instance with parameters for which Proposition 3(i) applies. In this case, a high absolute value of state induces sensor i to emphasize familiarity over quality such that it targets sensor l. In contrast, when the absolute value of state is low, quality matters more than familiarity, and i targets h. The reverse holds if case (ii) applies. When this intermediate region exits, then case (i) (i.e., high state favors high familiarity sensor) applies over this entire intermediate region if $r > \sqrt{w_{ih0}/w_{il0}}$, but case (ii) (i.e., high state favors high-quality sensor) applies over this entire intermediate region otherwise. Put together, these results show that state can play an instrumental role in mediating quality and familiarity even when target section is static (i.e., does not evolve over time). In the next section, we turn our attention to exploring how target selection evolves over time.

5.2. Target Evolution and Long Run Target Selection

Analogously to initial selection, when choosing its target in period t, any given sensor i needs to consider the quality of each sensor $j \in \mathcal{N} \setminus \{i\}$ as well as its current familiarity value $\psi_{ijt}(s_{t-1})$ for sensor j. What differs from the initial selection is that the familiarity function $\psi_{ijt}(s)$ may have evolved due to past targeting of j by i. As shown in Proposition 2, this familiarity function still depends on the initial belief variances $1/v_{ij0}$ and $1/w_{ij0}$ but it now also depends on (a) i's communication history with j as reflected by the number of times i targeted j in the past, and (b) the stickiness factor $f(q_i, q_j)$. In particular, $\psi_{ijt}(s_{t-1})$ strictly increases in the number of times i has already targeted j, and so the attractiveness of j as a future target for i increases every

[14] A complete closed-form analytical characterization of the region thresholds exists but it is algebraically cumbersome and not included for reasons of space.

time i targets j. This is because the signal received from j becomes more informative for i as its familiarity with j builds.

The stickiness factor $f(q_i, q_j)$ determines the gain in familiarity that results each time i targets j. It is strictly increasing in q_i and strictly decreasing in q_j; see (14). To understand these directional effects, first consider q_j, i.e., the quality of the targeted sensor. For the extreme case in which $q_j = \infty$, sensor j's state value (what it sends to sensor i) will be independent of its inference model parameters (see (4) and replace i with j), and so there is no information to be gained by i about j's model parameters. More generally, the higher the quality of sensor j, the less weight it places on its inference model in forming its state expectation, and so there is less for i to learn about j's model parameters as q_j increases. Therefore, i's gain in familiarity decreases in q_j. Next, consider q_i, i.e., the quality of the targeting sensor. As q_i increases, the targeting sensor's own state estimate becomes more precise. This, in turn, enables sensor i to better interpret and parse out the inference model information contained in what it receives from sensor j. To summarize, all else equal, (a) higher-quality sensors can build familiarity with other sensors more rapidly than can lower-quality sensors, and (b) lower-quality sensors result in larger familiarity gains if targeted.

To analyze how target selection evolves over time, it is helpful to introduce the following definition and result.

Definition 1 (**Dominance**). For two sensors $m, n \in \mathcal{N} \setminus \{i\}$, we say that m dominates n at time t from the perspective of sensor i (denoted by $m \succeq_{it} n$), if $Pr\big(\sigma_t^2(i, m, S_{t-1}) \leq \sigma_t^2(i, n, S_{t-1}) | \mathcal{H}_t\big) = 1$, where $\mathcal{H}_t$ is the history of communications up to time t ($\mathcal{H}_1 = \emptyset$).

In other words, $m \succeq_{it} n$ if sensor i almost surely prefers to target sensor m instead of n at time t given the history of all past communications. Using Definition 1, we can establish the following preservation result.

Lemma 1 **(Dominance Preservation)**. *If $m \succeq_{it} n$, then $m \succeq_{it'} n$, for all $t' > t$.*

This result establishes that dominance is preserved (i.e., persists) over time. Therefore, if some sensor n becomes dominated by some other sensor m from the perspective of i at some time t, then sensor n will never be targeted by i in the future. This allows sensor i to reduce its set of potential targets over time. This result enables us to analyze the long run communication network.

In what follows, we first consider two special cases, and then explore the general case.

Special Case 1 (Common Initial Familiarities that Vary by Sensor): Consider the case in which any given sensor i has a common initial familiarity with all other sensors $j \in \mathcal{N} \setminus \{i\}$, i.e., $v_{ij0} = v_{i0}$ and $w_{ij0} = w_{i0}$ for all j. This common initial familiarity can vary by sensor i. Let $h(i)$

denote the highest-quality sensor $j \in \mathcal{N} \setminus \{i\}$ from i's perspective. It follows from Proposition 3 that $h(i)$ dominates all other sensors (from the perspective of i) at time 1. Because dominance is preserved (Lemma 1), the communication network at any time (including the long run) is the same across all sample paths: regardless of state realizations, each sensor i always targets the highest quality sensor available $h(i)$. Put differently, the highest-quality sensor targets the second-highest quality sensor, and all other sensors target the highest quality sensor.

Special Case 2 (Equal Qualities with Initially More Familiar Sensors): Consider the case in which (a) the sensors are all of the same quality, and (b) for any given sensor i there exists some other sensor $\hat{j}(i)$ such that $v_{i\hat{j}0} \geq v_{ij0}$ and $w_{i\hat{j}0} \geq w_{ij0}$ for all $j \in \mathcal{N} \setminus \{i\}$. In other words, sensor i has higher initial familiarities with both of $\hat{j}(i)$'s inference model parameters than in any other sensor's parameters. It follows from part (iii) of Proposition 3 that $\hat{j}(i)$ dominates all other sensors (from the perspective of i) at time 1. Because dominance is preserved (Lemma 1), the communication network at any time (including the long run) is the same across all sample paths: regardless of state realizations, each sensor i always targets its initially most-familiar sensor $\hat{j}(i)$.

In general, however, sensors may differ in their qualities and any given sensor may have heterogenous familiarities with other sensors. In such a setting, an initially-dominant target (for any given sensor) may not exist. Therefore, we next develop results to help analyze this general case. To this end, let $\mathcal{S}_\infty \triangleq \{s_0, s_1, s_2, \cdots\}$ denote a long run sample path, i.e., the realization of states over an infinite horizon. Similarly, we denote by $\mathcal{S}_t \triangleq \{s_0, s_1, s_2, \cdots, s_t\}$ a finite sample path up to time t. We also let $\mathcal{S}'^t_\infty$ denote a sample path that is equivalent to $\mathcal{S}_\infty$ up to time t, but one which may deviate from $\mathcal{S}_\infty$ afterwards: $\mathcal{S}'^t_\infty \triangleq \mathcal{S}_t \cup \{s'_{t+1}, s'_{t+2}, \cdots\}$. To examine the long run networks that may arise, we first introduce the following definition.

Definition 2 (**Long Run Contacts**). Given a sample path $\mathcal{S}_\infty$, the set of long run contacts of sensor i is:

$$\mathcal{T}_i(\mathcal{S}_\infty) \triangleq \left\{ j \in \mathcal{N} \setminus \{i\} : \lim_{t \to \infty} \psi_{ijt}(s_{t-1}) = \infty | s_0, s_1, s_2, \cdots \in \mathcal{S}_\infty \right\}. \tag{20}$$

Remark 1 (**Infinitely-Often Communication**). It immediately follows from (15) that, along any sample path $\mathcal{S}_\infty$, sensor i targets sensor j infinitely often if, and only if, $j \in \mathcal{T}_i(\mathcal{S}_\infty)$.

If two (or more) sensors have the same quality, then depending on the initial familiarity of some sensor i with these equal-quality sensors, there might exist some sample paths along which the long run set of contacts of sensor i includes more than one sensor and sensor i keeps alternating between the sensors in its long run set of contacts such that it targets each of them infinitely often along the sample path. This alternating behavior is caused by the value of state in each period which, as noted earlier, plays a central role in target selection.

However, when qualities differ across sensors, we establish in what follows that for any given sensor i and along any fixed sample path $\mathcal{S}_\infty$: (a) $\mathcal{T}_i(\mathcal{S}_\infty)$ is a singleton, i.e., $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$, and (b) the unique long run contact in $\mathcal{T}_i(\mathcal{S}_\infty)$ can be identified in the almost sure sense in finite time, i.e., $\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S'}^{t^*}_\infty)$ *a.s.* for some $t^* < \infty$. These two results in turn will allow us to establish the following. At time zero, one can fully define the long run communication network as a *random directed graph*, i.e., a directed graph with given probabilities assigned to each link ij that indicate the probability that j will be the long run target for i. Furthermore, there exists a finite time after which the graph can be defined as a deterministic directed graph, i.e., with all probabilities being zero or one, that fully specifies the long run target for each sensor.

To establish these results, we start by presenting the following lemma.

LEMMA 2. *For any $\epsilon > 0$, there exists a fixed threshold $\bar{\psi}_\epsilon \in \mathbb{R}$ such that if $\psi_{ijt}(s_{t-1}) > \bar{\psi}_\epsilon$ and $\frac{q_j}{q_{j'}} > 1 + \epsilon$ then $j^*_{it} \neq j'$.*

The above lemma states that a sensor j' will not be targeted by a sensor i if (a) there is another sensor j of a higher quality than j', and (b) sensor i's familiarity with j reaches a fixed threshold. The importance of this result lies in the fact that the threshold is a fixed number, and hence, is independent of sensor i's familiarity with sensor j'. Thus, the above lemma holds regardless of how familiar i is with j' at time t: if i's familiarity with j passes the fixed threshold, then j' will not be targeted by i. This in turn allows us to show that, when sensor qualities are asymmetric (defined below), the set of long run contacts of each sensor i along any sample path $\mathcal{S}_\infty$ only includes one sensor.

DEFINITION 3 (**Asymmetric Qualities**). Sensor qualities are said to be asymmetric if, and only if, $q_j \neq q_{j'}$ for all $j, j' \in \mathcal{N}$ with $j \neq j'$.

PROPOSITION 4 **(Unique Long Run Contact)**. *If sensor qualities are asymmetric, then given any sample path $\mathcal{S}_\infty$, $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$ for all $i \in \mathcal{N}$.*

It is noteworthy that although the long run set of contacts of each sensor i has a unique member (when sensor qualities are asymmetric), this unique member is (a) sample path dependent, and (b) is not necessarily the highest quality sensor in $\mathcal{N} \setminus \{i\}$. As Proposition 3 showed, at $t = 1$ any given sensor i might target a sensor of lower quality than some other potential target. Due to the stickiness factor introduced in Proposition 2, this may create a momentum for sensor i to target the same sensor in future periods as well. This may result in a lower quality sensor dominating the higher quality sensor from the perspective of sensor i at some period t. Since dominance persists (see Lemma 1), the higher quality sensor may not be the long run contact of sensor i.

Using the above result, we next show that when the sensor qualities are asymmetric, the long run set of contacts of each sensor can be determined in finite time. That is, transient analysis is sufficient for characterizing the communication network that will be formed in the long run. This is because the role of state in target selection eventually vanishes, i.e., the effect of past targeting outweighs the role of state.

Proposition 5 **(Transient Analysis)**. *If sensor qualities are asymmetric, then along any sample path $\mathcal{S}_\infty$ there exists a finite period t^* such that for all $i \in \mathcal{N}$*

$$\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S}'^{t^*}_\infty) \qquad a.s.$$

The above results allow us to characterize the long run network of communications.[15] First, at time zero, this network can be viewed as a *random directed graph* $\vec{G}(\mathcal{N}, \mathcal{E}, \mathcal{P})$, where $\mathcal{N}$ (i.e., the set of sensors) is the set of vertices, $\mathcal{E} \triangleq \{(i,j) : i, j \in \mathcal{N}\}$ is the set of directed links, and $\mathcal{P}$ is a set of probability distributions that assign to link (i,j) probability p_{ij} defined as

$$p_{ij} \triangleq \sum_{\forall \mathcal{S}_\infty:\; j \in \mathcal{T}_i(\mathcal{S}_\infty)} Pr(\mathcal{S}_\infty). \tag{21}$$

Second, as the following result shows, the network can be defined as a deterministic directed graph after some finite time.

Proposition 6 **(Deterministic Random Directed Graph)**. *If sensor qualities are asymmetric, then there exists a finite time t^* such that given the sample path up to t^* (i.e., $\mathcal{S}_{t^*}$), the long run communication network can be defined as $\vec{G}(\mathcal{N}, \mathcal{E}, \mathcal{P})$ introduced above with the additional property that $p_{ij} \in \{0,1\}$ for all $i, j \in \mathcal{N}$.*

In Appendix A we show that all of the above results extend readily to a setting in which a sensor can simultaneously target $\lfloor k \rfloor > 1$ other sensors. The long run contact set of each sensor will, however, contain $min\{\lfloor k \rfloor, |\mathcal{N}| - 1\}$ sensors assuming that sensor qualities are asymmetric.

[15] The extension of the above results to settings with non-asymmetric sensors is straightforward, although as noted earlier, Proposition 4 may no longer hold.

5.3. Managerial Implications

As discussed in the introduction, firms in the process industries are increasingly using sensors to improve asset uptime and worker safety through better real-time knowledge of important operating-condition variables such as vibration (whose effective estimation is key to condition-based maintenance). Sensor communication enables improved estimation but it creates a dependency on other firms' sensors. That is, firm A's asset uptime depends to some degree on sensors from other firms if firm A's sensor targets these other sensors to improve its vibration estimate. These other firms are "information suppliers" to firm A. The results developed above allow a firm to predict which other firms will have a high likelihood of being its long-run information suppliers, both for (a) an individual asset on a particular worksite and (b) in aggregate across its fleet of assets over all worksites by examining the targeting frequency across its portfolio of sensors. Analogous to a manufacturer mapping its supplier base and monitoring corporate events (e.g., takeovers, divestments, bankruptcies, etc.) that might affect supply continuity, our results enable a firm to map its important sensor-based-information suppliers so that they can be monitored for events that might affect information continuity, e.g., a change in ownership of an important group of targeted sensors that might impact information sharing.

In addition to predicting and mapping its information suppliers, a firm might want to influence which suppliers are likely to be chosen (targeted) by it suppliers. For example, for reasons of corporate relationships or geopolitical considerations, a firm might be more comfortable being dependent on some firms rather than other firms. From our results above, the long-run probability p_{ij}, defined in (21), increases in v_{ij0} and w_{ij0}; that is, all else equal, sensor j is more likely to the long run contact of sensor i as i's initial familiarity (belief precision) in j increases. Therefore, (labeling firms in same manner as their sensors) if firm i would prefer to be dependent on firm j than k then firm i should invest upfront effort building familiarity with firm j's sensor, either strategically through corporate relationships or more tactically through personnel on a specific worksite. All else equal, firms that are more willing to promote initial familiarities with each others' sensors are more likely to end up as each others targets, and so firms in existing alliances may find that those relationships persist. Initial familiarities may also be enhanced through past exposure to firms in different environments (worksites). If firm i has experience targeting sensors from firm j in analogous environments previously then its initial familiarity with j in this new environment should be higher, making it more likely that j will be the long run contact. In this way, past targeting relationships between firms may promote future relationships and create opportunities for alliances that, although perhaps not originally planned, should be fostered because of mutual information dependence.

Looking at the quality lever, it follows from our earlier results that the probability that some sensor j will be the long run contact of some other sensor i increases in the quality of sensor j; that is, p_{ij}, defined in (21), increases in q_j. This then implies that firms deploying high-quality sensors are more likely to be the long-run contact of other firms' sensors and, therefore, to be more connected to other firms, i.e., be a central node. The implication of this observation is that firms with high-quality sensors will serve as an information supplier to many firms. Because, as discussed in the introduction, connections can build organizational relationships that require resources and attention, firms deploying high-quality sensors should anticipate the opportunities and consequences of being more central in IoT-enabled networks. Although beyond the scope of this particular research, this raises the interesting possibility of sensor-quality based pricing for information sharing to enable firms to monetize the value of having higher-quality sensors.

Firms may have less control over the underlying state dynamics that govern the environment's operating-condition variable. However, the communication pattern that emerges will be influenced by the underlying dynamics. This is illustrated in Figure 2 that presents the long run communication patterns for a collection of three sensors and nine different cases for the environment's AR(1) parameters (α, β). The sensor qualities and initial familiarities were chosen such that sensor 1 is the most attractive from a quality perspective but that sensor 3 is the most attractive from an initial familiarity perspective. Sensor 2 lies in between sensors 1 and 3 in that it represents mid values of both quality and initial familiarity. (Full details and discussion of this numeric study and others can be found in Unabridged Appendix I.)

Although no sensor is initially dominant from the perspective of any other sensor (by study design), observe in the top left panel of Figure 2 (low α and low β) that the long run contact of each sensor is always the high-quality sensor, i.e., 1 eventually always targets 2, and 2 and 3 eventually always target 1. However, observe in the bottom right panel (high α and high β) that the long run contact of each sensor is always the high-familiarity sensor, i.e., 1 and 2 target 3, and 3 targets 2. In contrast, at intermediate values of α and β, e.g., the center panel, the long run contact of each sensor typically depends on the sample path: for some sample paths, the higher quality sensor wins and for others the higher-familiarity sensor wins. It is important to emphasize that although the underlying state dynamics affect the communication pattern, a firm can (as discussed above) influence its centrality in the pattern irrespective of the underlying dynamics through its sensor

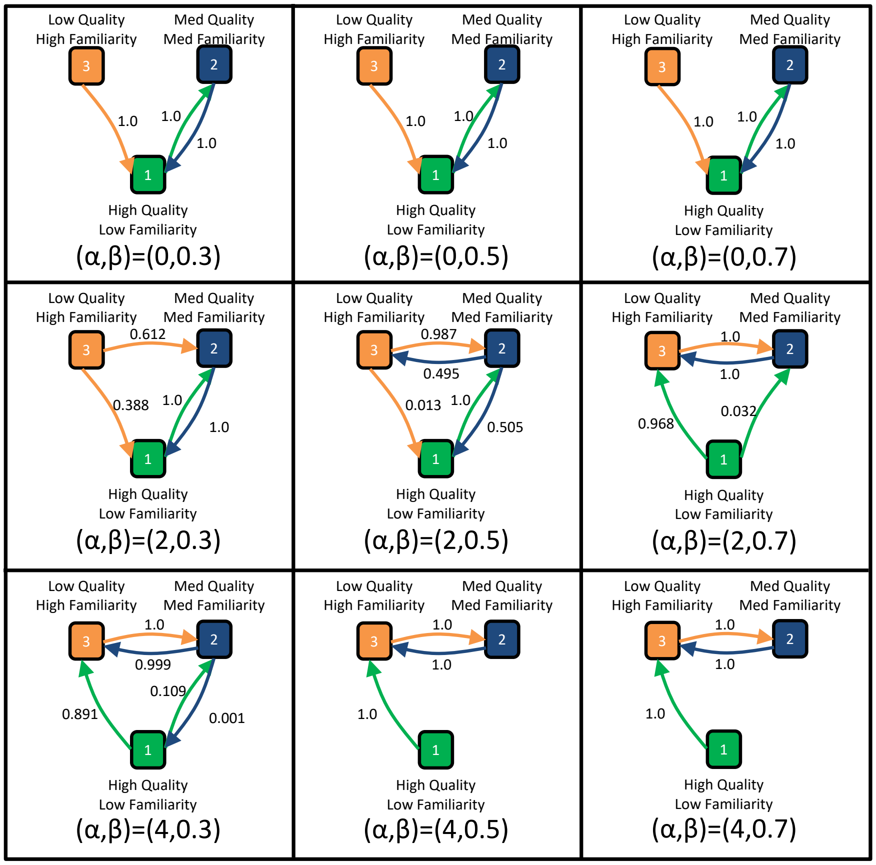

Figure 2 The influence of underlying state dynamics on communication pattern.

quality and it willingness (or not) to increase others' initial familiarity with its sensor model. It can also influence who it is likely to target (irrespective of the underlying state dynamics) by increasing its familiarity with targets it sees as desirable for certain corporate reasons.

6. Sensors of Unknown Qualities

To this point, we have assumed that sensor qualities are common knowledge. That might not always be the case; a sensor deployed by one firm may have only limited knowledge about the quality of a sensor deployed by a different firm.

In what follows, we allow sensor qualities to be ambiguous to other sensors. We do so by assuming that any given sensor i believes the quality of sensor $j \in \mathcal{N} \setminus \{i\}$ is contained in a set of possible values (which we refer to as the ambiguity set) with each possible value having some associated probability. We adopt a robust optimization framework in which sensors select their target so as to be robust to this ambiguity, while trying to achieve the best possible improvement in their estimation as in the previous sections. To this end, let $\mathbb{P}_i^Q$ denote the joint probability that sensor i assigns to the possible qualities of all other sensors. We consider a robust version of the target selection problem (6), where similar to the previous sections we assume $\lfloor k \rfloor = 1$ for expositional ease. In each time period t, any given sensor i follows a targeting strategy that is defined by a $|\mathcal{N}| - 1$ dimensional probability vector with elements representing the probability that sensor i targets sensor $j \in \mathcal{N} \setminus \{i\}$. In particular, we assume sensor i's problem at time t is to find the targeting strategy:

$$\pi^*_{it} = \arg \inf_{\pi \in \Pi_i} y^\pi_{it} \tag{22}$$

where

$$y^\pi_{it} = \inf_{y_\epsilon \in \mathbb{R}_+} y_\epsilon \tag{23}$$

$s.t.$

$$\mathbb{P}^Q_i \Big\{ \min_{\tilde{s}_{it} \in \mathbb{R}} \mathbb{E}_{\pi, S_t \sim F^\pi_{it}} \big[\tilde{s}_{it} - S_t \big]^2 \le y_\epsilon \Big\} \ge 1 - \epsilon. \tag{24}$$

That is, at time t each sensor i optimizes over the set of targeting strategies Π_i (which contains all deterministic and/or randomized strategies) to find the current targeting strategy that minimizes y^π_{it}, where y^π_{it} represents the robust "cost" of a targeting strategy π. This cost is defined as the $(1-\epsilon)$-percentile (with respect to $\mathbb{P}^Q_i$) of the sensor i's estimation squared error if it follows targeting strategy π. In (24), F^π_{it} is the posterior distribution of sensor i's belief about the state after applying the targeting strategy π at time t.[16] Parameter $\epsilon \in [0, 1]$ represents the level of *optimism*, where $\epsilon = 0$ yields robust optimization with respect to the worst-case (a pessimistic scenario), and $\epsilon = 1$ yields robust optimization with respect to the best-case (an optimistic scenario).

The optimal targeting strategy given by (22) need not be deterministic in general: a randomized strategy might outperform any deterministic strategy due to the chance constrained optimization nature of problem (22)-(24). This is formalized in Lemma EC.5 in Appendix H. As we will see in our numerical studies below, this randomization may cause a sensor to have a long run set of contacts that includes more than one member even if qualities are asymmetric, which is in stark contrast to the singleton result we established when qualities are known (Proposition 4). That is, in order to be robust to the fact that qualities of other sensors are not perfectly known, each sensor may end up building enough familiarity with more than one other sensor in the long run, and go back and forth between them infinitely often along any given sample path. Even in that case, the long-run contact set, though not a singleton, can be stochastically predicted, and so our earlier managerial implications about information-supplier mapping still apply. However, there are

[16] This posterior distribution depends on the element of i's ambiguity set (i.e., the particular q_j values) as well as the past targeting history (through the familiarity function at time t which in turn depends on q_j values), but these dependencies are suppressed for ease of notation.

conditions under which one can restrict attention to the set of deterministic policies without any loss. Below, we present on such sufficient condition.

Proposition 7 **(Deterministic Communication)**. *Suppose that at period t for all $j \in \mathcal{N} \setminus \{i\}$ we have $V_{\tilde{j}} \leq_{s.t.} V_j$ for some $\tilde{j} \in \mathcal{N} \setminus \{i\}$. Then π^*_{it} defined in (22) prescribes that sensor i targets sensor $\tilde{j}$ at period t almost surely regardless of ϵ.*

Proposition 7 establishes a connection between cases with unknown qualities and those with known qualities. When qualities are known, sensor i targets the sensor that provides the lowest signal variance, i.e., the most informative signal. When qualities are unknown, this deterministic comparison has a stochastic counterpart: if the signal variance from one sensor $\tilde{j}$ is *stochastically* lower than that of other sensors, sensor i targets sensor $\tilde{j}$ with probability one, regardless of the optimism level, ϵ. Thus, sensor i still behaves deterministically for any robustness level imposed by ϵ. However, this deterministic behavior may not hold if sensor i assigns probabilities to unknown qualities in a way that no one sensor's signal variance stochastically dominates the others.

We numerically explored how the tradeoffs in the known-quality setting between (a) quality, (b) familiarity, and (c) state can be affected by the underlying ambiguity around qualities and/or by the level of optimism of sensors. We briefly summarize our observations here and refer the reader to Appendix I (Unabridged Version) for full details and discussion. For the case of known qualities, we analytically established earlier that target selection is deterministic (i.e., sample path independent) and time-invariant in the case of common initial familiarities (see earlier Special Case 1): the highest-quality sensor targets the second-highest quality sensor, all other sensors target the highest quality sensor. However, when qualities are unknown, the long-run contact might not be deterministic even under a deterministic targeting policy each period; see Study 3 in Appendix I. We also observed (Study 4 in Appendix I) that the optimism parameter ϵ strongly influences long-run target selection. Sensors are more pessimistic (optimistic) about potential sensor qualities when ϵ is low (high) and this in turn influences the emphasis placed on familiarity versus possible qualities, which in turn influences the role of state in target selection. The optimism parameter used by a sensor will depend on the firm that deployed it. Thus, organizational attitudes towards ambiguity will impact target selection and the resulting communication network that evolves over time.

7. Extensions

In what follows we summarize a number of extensions of our base model. Full details and discussion of these extensions can be found in the appendices.

We assumed that each sensor could target only one sensor ($\lfloor k \rfloor = 1$) in each period, both for expositional ease and because that is when the targeting choice is most crucial. However, we relax this assumption in Appendix A and extend our results to a general channel capacity $\lfloor k \rfloor > 1$. We establish that each sensor will follow a rank-ordering structure in its targeting: in each period it will (a) order all the sensors from lowest to highest based on the variance criterion established earlier, i.e., (10), and then (b) pick the $\lfloor k \rfloor$ best-ranked, i.e., lowest-variance, sensors as its targets. We prove that our key results (Propositions 1, Propositions 2, Lemma 2, and Proposition 4) extend in a natural fashion based on this rank-ordering structure. The extension of other results is then straightforward. Our results also carry over directly to a setting in which the channel capacity $\lfloor k \rfloor$ varies by sensor.

We assumed that the initial belief of a sensor about another sensor's inference model parameters are accurate in their means. The extension to a setting in which the means of these normal distributions are not correct is contained in Appendix B which establishes that the key results in the paper still hold.

We assumed that each sensor has a slower-but-more-accurate estimation approach in place (e.g., technician inspection or remote offloading) that has a delay of one period and is perfectly accurate. In Appendix C, we analyze various relaxations in which the slower-but-more-accurate approach has a general delay and its estimation may or may not be perfectly accurate. For a general delay with perfect accuracy, i.e., when state realization is simply delayed, we show that the following key results still hold: if sensor qualities are asymmetric then (a) there exists a unique, albeit sample-path dependent, long run contact; (b) this contact can be determined in finite time; and (c) there exists a finite time after which the long run communication network can be defined as a deterministic directed graph. That is, Propositions 4, 5, and 6 all continue to hold; see Propositions EC.4, EC.5, and EC.6 in Appendix C for details. We then establish that these results continue to hold under a general delay with imperfect accuracy, i.e., when the slower-but-more-accurate approach does not reveal state perfectly to the sensor. We also explore settings in which the slower-but-more-accurate approach does not exist and so there is no state realization except for the first period or not even in the first period.

In Appendix D, we extend our base model to the case where the error terms in sensor readings are correlated. We first consider a "symmetric correlated measurement errors" case in which the correlation between sensor errors is the same for all pairs of sensors. We establish that the optimal target choice will be the same as when there is no correlation and, therefore, our key results developed in the main body still hold and the network that is formed among sensors (both in the short and long-term) will be the same as the one without correlated errors. We then consider the "asymmetric correlated measurement errors" case in which the correlation in measurement errors between a sensor i and a sensor j is allowed to depend on i and j. We show that this induces a target selection criterion that is not in general equivalent to the one established in the main body. Hence, a sensor might target a different sensor than it would if there was no correlation. However, we are able to establish a condition (related to the ratio of correlation difference to sensor qualities) under which the target selection criterion for each sensor will be equivalent to that under no correlation and hence our key results will hold. Taking both cases together, it follows that the results for the case of uncorrelated errors can be quite informative of what happens even if errors are correlated.

The extension to the setting in which sensors care about the quality of future estimates and not just of the current-period estimate is analyzed in Appendix E. We establish that this non-myopic problem lies in the general class of restless multi-armed bandit problems. Nonetheless, we show that if the discount factor (for future estimates) is below a certain threshold then it is optimal for sensors to act myopically (similar to the setting we studied earlier). For any arbitrary discount factor, we also establish a sufficient condition for myopic targeting to be optimal in a given state.

Finally, in Appendix F, we allow sensors to update their own inference parameters in a Bayesian fashion. We formally establish that if each sensor's initial precision of its own inference model parameters exceeds a certain threshold (and this is publicly known) then one's own inference parameter updating will not alter the target selection of any of the sensors. This implies that if pre-deployment training data sets are large enough to result in sufficiently high initial precisions then network formation among sensors can be accurately studied without assuming that sensors update their own inference models after deployment. Since in practice firms typically use large amount of training data pre-deployment, minor updating of one's own inference model parameters can be safely ignored (thus, the setting we studied earlier). However, because different firms may use different sensor technologies, proprietary algorithms, and different training data sets pre-deployment, their inference parameters will be private and therefore learning about other sensor's parameters is still fruitful.

8. Conclusions

Much of the promise of the IoT stems from the idea that better operational decisions will be enabled by a vast array of sensors that provide almost real-time knowledge of the state of things. In the process industries, assets and personnel from different firms are often located in close proximity on the same worksite, and sensors are now being widely deployed to monitor key operating-condition variables (vibration, for example) in an effort to improve asset uptime and worker safety through more-informed condition-based maintenance. Sensors are not perfectly accurate but the sharing of estimates across sensors can help improve estimation. However, sensors cannot solicit information from all other sensors in their environment due to very real communication constraints, and sensors may not have full knowledge of the inference models used by sensors deployed by other firms.

We characterize the initial and long run communication network—who talks to who—for an arbitrary collection of sensors that do not know each others' underlying inference models and that may not know each others' qualities. We establish that the state of the environment plays a key role in determining the weights placed on quality and familiarity (knowledge of another's inference model) when selecting a target. We establish that if sensors differ in their qualities then each sensor will eventually target a single sensor in all future periods. This long run target, however, can vary by sensor and is sample-path dependent because state values influence the weight sensors put on familiarity versus quality. We establish that the long run communication network that forms between sensors can be fully defined at time zero as a random directed graph and that one can probabilistically predict the long-run communication patterns that will emerge. When qualities are not common knowledge, we show that a firm's ambiguity attitude can play an important role in target selection. Our work sheds light on what kind of communication networks develop over time, and this enables managers to not only make predictions about which other firms their sensors will interact with but to also influence the communication outcomes through the levers of sensor qualities and initial familiarities. This predictive ability and managerial control is important in light of the fact that sensor communications build organizational ties that require attention and resources. For example, our work enables a process-industry firm to map the sensor-based-information suppliers that it will depend on, at least to some extent, for improved uptime performance.

This specific research could be extended in a number of directions. The sensors might not operate in the same environment but instead operate in correlated environments such that signals are still somewhat informative to each other. The environment might evolve according to a more general model than the AR (1) model we used to generate insights. We assumed that the receiver is always able to respond to the sender (that is, when i communicates to j, j responds to i) and leave it to future research to explore potential relaxations. We assumed time-invariant sensor qualities. If sensor qualities degrade significantly over time then this would introduce an important time dependence between quality and time that would cause the familiarity function dynamics (Proposition 2) to depend not only on the number of times a sensor was targeted but when it was targeted.

Because the goal of this research was not to examine a highly specific application but rather to establish and analyze a general information-quality based communication framework that applies to a broad range of emerging IoT settings, we were intentionally silent about (a) the actions of the sensor-owning entities and (b) the incentives of these entities to share information. With regard to (a), because sensor targeting in our framework faces a constraint on the number of sensors targeted in a period, we could, without loss of generality, ignore any actions taken based on sensor estimates under the mild assumption that better estimation allows better actions (and hence, sensors have the objective of providing their entities with the best estimates). However, if there were financial costs to soliciting estimates from other sensors, then the value of the improved estimation would need to be taken into account when targeting, and that would require a model of how optimal actions (e.g., condition-based maintenance) and pay-offs (e.g., asset uptime) depend on estimation performance. We leave the development of models that consider targeting costs and the value of better estimation for future research. With regard to (b), we assumed a particular partial-information sharing regime in which the entities only share state estimates but not their sensors' proprietary inner processes (e.g., inference models or readings). Future research could study decisions regarding whether or not to share (and what information to share) under different incentive structures.

Finally, we have focused on the information quality motive for sensor communication but firms are also interested in information completeness in which the states of distinct elements are combined to provide an overall system state. In general, the IoT presents many opportunities to explore how to improve and exploit information quality and information completeness in various operations-related domains. We hope our paper motivates further research in this area.

References

Acemoglu, Daron, Kostas Bimpikis, Asuman Ozdaglar. 2014a. Dynamics of information exchange in endogenous social networks. *Theoretical Economics* **9**(1) 41–97.

Acemoglu, Daron, Mohamed Mostagir, Asuman Ozdaglar. 2014b. State-dependent opinion dynamics. *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4773–4777.

Ashton, Alison Hubbard, Robert H Ashton. 1985. Aggregating subjective forecasts: Some empirical results. *Management Science* **31**(12) 1499–1508.

Atanasov, Pavel, Phillip Rescober, Eric Stone, Samuel A Swift, Emile Servan-Schreiber, Philip Tetlock, Lyle Ungar, Barbara Mellers. 2016. Distilling the wisdom of crowds: Prediction markets vs. prediction polls. *Management Science* **63**(3) 691–706.

Ayaz, Emine. 2014. Autoregressive modeling approach of vibration data for bearing fault diagnosis in electric motors. *Journal of Vibroengineering* **16**(5) 2130–2138.

Ballotta, Luca, Luca Schenato, Luca Carlone. 2019. Computation-communication trade-offs and sensor selection in real-time estimation for processing networks. *arXiv* arXiv–1911.

Bassamboo, Achal, Ruomeng Cui, Antonio Moreno. 2018. The wisdom of crowds in operations: Forecasting using prediction markets. *Working Paper, Available at SSRN* `https://ssrn.com/abstract=2679663`.

Bates, John M, Clive WJ Granger. 1969. The combination of forecasts. *Journal of the Operational Research Society* **20**(4) 451–468.

BehrTech. 2020. 5 iot applications for offshore monitoring in oil and gas. Available at `https://behrtech.com/blog/5-iot-applications-for-offshore-monitoring-in-oil-and-gas/` Accessed May 21, 2020.

Bren, Austin, Soroush Saghafian. 2019. Data-driven percentile optimization for multi-class queueing systems with model ambiguity: theory and application. *INFORMS Journal on Optimization* **1**(4) 267–287.

Chen, Cheng, Xue-Ling Zhao, Zhan-Hong Li, Zhi-Gang Zhu, Shao-Hong Qian, Andrew J Flewitt. 2017. Current and emerging technology for continuous glucose monitoring. *Sensors* **17**(1) 182.

Delage, Erick, Dan A Iancu. 2015. Robust multistage decision making. *The Operations Research Revolution*. INFORMS, 20–46.

Delage, Erick, Shie Mannor. 2010. Percentile optimization for Markov decision processes with parameter uncertainty. *Operations Research* **58**(1) 203–213.

Evans, Peter, Marco Annunziata. 2012. Industrial internet: Pushing the boundaries of minds and machines. *General Electric Report* .

Fisher, Marshall, Ananth Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research* **44**(1) 87–99.

FlukeCorp. 2020. An introduction to machinery vibration. Available at `https://www.reliableplant.com/Read/24117/introduction-machinery-vibration` Accessed May 21, 2020.

Gaur, Vishal, Saravanan Kesavan, Ananth Raman, Marshall L Fisher. 2007. Estimating demand uncertainty using judgmental forecasts. *Manufacturing & Service Operations Management* **9**(4) 480–491.

GEPower. 2018. Electrical rotating machine apm overview. Available at `https://www.gepowerconversion.com/sites/default/files/GEA33604_20Motor_Fleet_APM_White_Paper_March_2019.pdf`.

Hall, David L, James Llinas. 1997. An introduction to multisensor data fusion. *Proceedings of the IEEE* **85**(1) 6–23.

Han, Duo, Junfeng Wu, Huanshui Zhang, Ling Shi. 2017. Optimal sensor scheduling for multiple linear dynamical systems. *Automatica* **75** 260–270.

IEC. 2015. *Internet of Things: Wireless Sensor Networks*. International Electrotechnical Commission.

IKM. 2019. Condition monitoring: Oil & gas. Available at `https://www.ikm.com/getfile.php/1343855-1559300145/IKM20Selskaper/IKM20Instrutek/Nedlastninger/Datablad20og20brosjyrer/Tilstandskontroll/CM20Offshore_low.pdf`.

Jernigan, Stephanie, David Kiron, Sam Ransbotham. 2016. Data sharing and analytics are driving success with IoT. *MIT Sloan Management Review* **58**(1).

Khaleghi, Bahador, Alaa Khamis, Fakhreddine O Karray, Saiedeh N Razavi. 2013. Multisensor data fusion: A review of the state-of-the-art. *Information Fusion* **14**(1) 28–44.

Leal, Yenny, Winston Garcia-Gabin, Jorge Bondia, Eduardo Esteve, Wifredo Ricart, Jose-Manuel Fernández-Real, Josep Vehí. 2010. Real-time glucose estimation algorithm for continuous glucose monitoring using autoregressive models. *Journal of diabetes science and technology* **4**(2) 391–403.

Lennon Olsen, Tava, Brian Tomlin. 2019. Industry 4.0: Opportunities and challenges for operations management. *Manufacturing & Service Operations Management (forthcoming)* .

Liyanage, Liwan H, J George Shanthikumar. 2005. A practical inventory control policy using operational statistics. *Operations Research Letters* **33**(4) 341–348.

Marr, Bernard. 2016. What everyone must know about industry 4.0. Available at `https://www.forbes.com/sites/bernardmarr/2016/06/20/what-everyone-must-know-about-industry-4-0`.

McKinsey. 2015. *The Internet of Things: Mapping the value beyond the hype*. McKinsey Global Institute.

Mišić, Velibor, Georgia Perakis. 2019. Data analytics in operations management: A review. *Manufacturing & Services Operations Management (forthcoming)* .

Mitchell, Harvey B. 2007. *Multi-sensor data fusion: An introduction*. Springer Science & Business Media.

Momot, Ruslan, Elena Belavina, Karan Girotra. 2019. The use and value of social information in selective selling of exclusive products. *Management Science (forthcoming)* .

Nemirovski, Arkadi, Alexander Shapiro. 2006. Convex approximations of chance constrained programs. *SIAM Journal on Optimization* **17**(4) 969–996.

Perakis, Georgia, Guillaume Roels. 2008. Regret in the newsvendor model with partial information. *Operations Research* **56**(1) 188–203.

Ran, Xukan, Haoliang Chen, Zhenming Liu, Jiasi Chen. 2017. Delivering deep learning to mobile devices via offloading. *Proceedings of the Workshop on Virtual Reality and Augmented Reality Network*. 42–47.

Saghafian, Soroush, Brian Tomlin. 2016. The newsvendor under demand ambiguity: Combining data with moment and tail information. *Operations Research* **64**(1) 167–185.

Sethi, Rajiv, Muhamet Yildiz. 2016. Communication with unknown perspectives. *Econometrica* **84**(6) 2029–2069.

Shi, Ling, Huanshui Zhang. 2012. Scheduling two gauss–markov systems: An optimal solution for remote state estimation under bandwidth constraint. *IEEE Transactions on Signal Processing* **60**(4) 2038–2042.

Simchi-Levi, David. 2010. *Operations Rules: Delivering customer value through flexible operations*. MIT Press.

Sparacino, Giovanni, Francesca Zanderigo, Stefano Corazza, Alberto Maran, Andrea Facchinetti, Claudio Cobelli. 2007. Glucose concentration can be predicted ahead in time from continuous glucose monitoring sensor time-series. *IEEE Transactions on biomedical engineering* **54**(5) 931–937.

Swaminathan, Jayashankar M, Sridhar R Tayur. 2003. Models for supply chains in e-business. *Management Science* **49**(10) 1387–1406.

Syed, SA, M Pai. 2016. Antifriction bearing diagnostics in a manufacturing industry: a case study. *J Mech Eng Autom* **6**(5A) 58–62.

Thanagasundram, Suguna, Fernando Soares Schlindwein. 2006. Autoregressive order selection for rotating machinery. *International Journal of Acoustics and Vibration* **11**(3) 144–154.

Tong, Jordan, Daniel Feiler. 2017. A behavioral model of forecasting: Naive statistics on mental samples. *Management Science* **63**(11) 3609–3627.

Tsoukalas, Gerry, Brett Hemenway Falk. 2019. Token-weighted crowdsourcing. *Management Science (forthcoming)* .

Vitus, Michael P, Wei Zhang, Alessandro Abate, Jianghai Hu, Claire J Tomlin. 2012. On efficient sensor scheduling for linear dynamical systems. *Automatica* **48**(10) 2482–2493.

Wallis, Kenneth F. 2011. Combining forecasts–forty years later. *Applied Financial Economics* **21**(1-2) 33–41.

Winkler, Robert L, Robert T Clemen. 2004. Multiple experts vs. multiple methods: Combining correlation assessments. *Decision Analysis* **1**(3) 167–176.

Winnig, Laura. 2016. Ge's big bet on data and analytics. *MIT Sloan Management Review* **57**(3).

Wu, Shuang, Kemi Ding, Peng Cheng, Ling Shi. 2020. Optimal scheduling of multiple sensors over lossy and bandwidth limited channels. *IEEE Transactions on Control of Network Systems* .

Xiong, Feiyu, Brian R Hipszer, Jeffrey Joseph, Moshe Kam. 2011. Improved blood glucose estimation through multi-sensor fusion. *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*. IEEE, 377–380.

Xu, Dianlei, Tong Li, Yong Li, Xiang Su, Sasu Tarkoma, Pan Hui. 2020. A survey on edge intelligence. *arXiv preprint arXiv:2003.12172* .

Yang, Wen, Guanrong Chen, Xiaofan Wang, Ling Shi. 2014. Stochastic sensor activation for distributed state estimation over a sensor network. *Automatica* **50**(8) 2070–2076.

Appendix A: General Channel Capacity

In our base model we assumed that each sensor can only target one sensor ($\lfloor k \rfloor = 1$) in each period although its target may vary across periods. We focused on this case because $\lfloor k \rfloor = 1$ represents the setting in which choosing the right target is the most crucial. We now extend our results by allowing a general value for $\lfloor k \rfloor$. As we will show, there is a *rank-ordering structure*: each sensor in each period selects the $\lfloor k \rfloor$ other sensors that provide it with the lowest variance of signal (about the state) in that period and targets them. When the channel capacity weakly exceeds the number of all other sensors in the environment ($\lfloor k \rfloor \geq |\mathcal{N}| - 1$), the results become trivial because each sensor will target all other sensors. Thus, without loss of generality, in what follows we assume ($\lfloor k \rfloor < |\mathcal{N}| - 1$). The rank-ordering structure implies (proven below) that each sensor's set of long run contacts have exactly $\lfloor k \rfloor$ members assuming that sensor qualities are asymmetric (see Definition 3 in the main body). We focus on extending some of our main results (Propositions 1, Propositions 2, Lemma 2, and Proposition 4) with the understanding that the extension of other results become straightforward given these new extensions.

We start by providing a generalization for Proposition 1 presented in the main body.

Proposition EC.1 **(Target Selection and Variance Reduction)**. *Suppose in each period sensor i is allowed to communicate to a general number of other sensors denoted by $\lfloor k \rfloor$. Then $a^*_{ijt} = \mathbb{1}_{\{j \in \mathcal{J}^*_{it}\}}$, where $\mathbb{1}_{\{\cdot\}}$ is the indicator function and $\mathcal{J}^*_{it} \subseteq \mathcal{N} \setminus \{i\}$ is the set of all sensors j that have the $\lfloor k \rfloor$ lowest value of*

$$\sigma_t^2(i, j, s_{t-1}) = \frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4}. \tag{EC.1}$$

Proof of Proposition EC.1: Proofs of all results in the Appendices are in Appendix H. □

We next present the following result, which shows that Proposition 2 still holds when the number of targets in each period is allowed to be any general number.

Proposition EC.2 **(Familiarity Dynamics)**. *Suppose in each period sensor i is allowed to communicate to a general number of other sensors denoted by $\lfloor k \rfloor$. Then for any $s \in \mathbb{R}$:*
(i) $\psi_{ij,t+1}(s) = \psi_{ijt}(s) + \delta(q_i, q_j, a_{ijt})$, where $\delta(q_i, q_j, a_{ijt}) \triangleq f(q_i, q_j) a_{ijt}$ and

$$f(q_i, q_j) \triangleq \frac{(1 + q_i^2)}{q_j^2(1 + q_i^2 + q_j^2)}. \tag{EC.2}$$

(ii) For all $t = 1, 2, 3, \cdots$, we have

$$\psi_{ij,t+1}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j) \sum_{l=1}^{t} a_{ijl}. \tag{EC.3}$$

Next, we extend Lemma 2. The following result shows that if there are $\lfloor k \rfloor$ other sensors that (a) have a higher quality than a given sensor j', and (b) the familiarity of sensor i with them is above a threshold in a period t, then sensor i will not target sensor j in that period. It then follows from the results established in the paper that all such $\lfloor k \rfloor$ sensors dominate sensor j' from the perspective of sensor i in period t. Due to the dominance preservation results established in the paper, this means that sensor i will not target sensor j' in any future period either. This, in turn, allows us to establish that the long-run set of contacts of sensor i will have exactly $\lfloor k \rfloor$ members. Notably, this means that each sensor will eventually settle on $\lfloor k \rfloor$ other sensors, and will not change the sensors it targets in future periods. This long-run set of contacts, however, depends on the sample path, and might differ across different realizations of state over time.

Lemma EC.1. *Suppose in each period sensor i is allowed to communicate to a general number of other sensors denoted by $\lfloor k \rfloor$. Consider a subset of sensors $\mathcal{J} \subseteq \mathcal{N} \setminus \{i\}$ satisfying $|\mathcal{J}| = \lfloor k \rfloor$. For any $\epsilon > 0$, there exists a fixed threshold $\bar{\psi}_\epsilon \in \mathbb{R}$ such that if for all $j \in \mathcal{J}$ (a) $\psi_{ijt}(s_{t-1}) > \bar{\psi}_\epsilon$, and (b) $\frac{q_j}{q_{j'}} > 1 + \epsilon$, then $a^*_{ij't} = 0$.*

Proposition EC.3 **(Set of Long Run Contacts)**. *Suppose in each period sensor i is allowed to communicate to a general number of other sensors denoted by $\lfloor k \rfloor$. If sensor qualities are asymmetric, then given any sample path $\mathcal{S}_\infty$, $|\mathcal{T}_i(\mathcal{S}_\infty)| = \lfloor k \rfloor$ for all $i \in \mathcal{N}$.*

Appendix B: Erroneous Initial Beliefs about Means

In our base model we assumed that at time $t = 0$ sensor i believes that sensor j's inference model parameters come from normal distributions that have correctly calibrated means ($N(\hat{\alpha}_j, 1/v^2_{ij0})$ and $N(\hat{\beta}_j, 1/w^2_{ij0})$). Since in practice, inference models are carefully built based on training data sets, it is not a strong assumption to assume that means of these distributions are relatively accurate. However, we now extend our analyses and consider the setting in which these means are erroneous as well.

In particular, suppose that at time $t = 0$, sensor i believes that sensor j's inference model parameters come from normal distributions $N(\tilde{\alpha}_j, 1/v^2_{ij0})$ and $N(\tilde{\beta}_j, 1/w^2_{ij0})$, where $(\tilde{\alpha}_j, \tilde{\beta}_j) \neq (\hat{\alpha}_j, \hat{\beta}_j)$. Let $\tilde{\boldsymbol{\mu}} = (\tilde{\alpha}_j, \tilde{\beta}_j)$, $\boldsymbol{\mu} = (\hat{\alpha}_j, \hat{\beta}_j)$, and denote the sensor i's belief about sensor j's inference model parameters at time t by $\mathbf{X} = (\alpha_{ijt}, \beta_{ijt})^T$ (a column vector). Suppose sensor i communicates to sensor j at time $t = 1$. It follows from the general results provided under Method 2 of proof of Proposition 2 that upon receiving the signal

$$\mathbf{Y} = \alpha_{ij1} + \beta_{ij1} s_0 + q_j^2 (S_1 + \epsilon_{j1})$$

from sensor j, sensor i's posterior joint distribution on sensor i's parameters is

$$\mathbf{X}|\mathbf{Y} \sim N(\boldsymbol{\Sigma}[\mathbf{A}^T \mathbf{L}(\mathbf{Y} - \mathbf{b}) + \boldsymbol{\Lambda}\tilde{\boldsymbol{\mu}}], \boldsymbol{\Sigma}), \tag{EC.4}$$

where $\mathbf{A} = (1, s_0)$, $\mathbf{b} = \mathbb{E}\left[q_j^2(S_1 + \epsilon_{j1})\right]$, $\boldsymbol{\Lambda} = \boldsymbol{\Sigma}^{-1}_{ij0}$, $\mathbf{L} = f(q_i, q_j)$, and

$$\boldsymbol{\Sigma} = \left[\boldsymbol{\Lambda} + \mathbf{A}^T \mathbf{L} \mathbf{A}\right]^{-1}. \tag{EC.5}$$

The difference between the posterior joint distribution obtained in (EC.4) and what one would get under the assumption made in the main body (i.e., correct means at times zero) is that in the latter case $\tilde{\boldsymbol{\mu}}$ would be replaced with $\boldsymbol{\mu}$. However, it follows from the proof of Proposition 2, that the results in Proposition 2 (and hence, other results in the paper) hold. Furthermore, under some conditions, one can also show that the initial difference between starting with mean beliefs $\tilde{\boldsymbol{\mu}}$ instead of $\boldsymbol{\mu}$ disappears as $t \to \infty$ assuming i targets j infinitely often).

Appendix C: Delayed, Imperfect or No State Observations

In our base model we assumed that each sensor has a slower-but-more-accurate estimation approach in place (e.g., technician inspection or remote offloading) that has a delay of one period and is perfectly accurate. We analyze various relaxations in which the slower-but-more-accurate approach has a general delay but is perfectly accurate (Case I) or has a general delay with imperfect observations (Case II). We also examine settings in which the only state realization is at time 0 (see Cases III and IV).

Case I: Delayed-but-Perfect State Observations

Recall that each sensor i has its own model of state dynamics given by (3):

$$S_{it} = \hat{\alpha}_i + \hat{\beta}_i S_{t-1} + \tilde{\epsilon}_t, \tag{EC.6}$$

where $\hat{\alpha}_i$ and $\hat{\beta}_i$ are sensor i's estimates of the process parameters α and β. In the main body, we assumed that right before period t starts, the realized value of the previous period state (s_{t-1}) becomes publicly known (i.e., is revealed to each sensor), and the system moves to the next period. Here, we relax this assumption and assume that there is a delay of n periods in observations made about the sate. That is, we now assume that right before period t starts, the value of S_{t-n} (for some positive integer $n \le t$) is observed perfectly by all sensors. Note that this is a general extension and incorporates various scenarios (including the one analyzed in the main body) as a special case. For example, as we will discuss later in this appendix, setting $n = t$ means that no state observation can be made beyond that made at the very beginning of the horizon (i.e., s_0).

Observe that replacing $S_{i,t-1} = \hat{\alpha}_i + \hat{\beta}_i S_{t-2} + \tilde{\epsilon}_{t-1}$ in (3) and repeating this process yields:

$$S_{it} = \hat{\alpha}_i \sum_{k=0}^{n-1} \hat{\beta}_i^k + \hat{\beta}_i^n S_{t-n} + \sum_{k=0}^{n-1} \hat{\beta}_i^k \tilde{\epsilon}_{t-k}. \tag{EC.7}$$

Hence, we have

$$\begin{aligned}\mathbb{E}[S_{it}|S_{i,t-n} = s_{i,t-n}] &= \hat{\alpha}_i \sum_{k=0}^{n-1} \hat{\beta}_i^k + \hat{\beta}_i^n s_{t-n} \\ &= \hat{\alpha}_i \frac{1-\hat{\beta}_i^n}{1-\hat{\beta}_i} + \hat{\beta}_i^n s_{t-n}\end{aligned} \tag{EC.8}$$

and

$$\begin{aligned}Var[S_{it}|S_{i,t-n} = s_{i,t-n}] &= \sum_{k=0}^{n-1} \hat{\beta}_i^{2k} Var[\tilde{\epsilon}_{t-k}] \\ &= \frac{1-\hat{\beta}_i^{2n}}{1-\hat{\beta}_i^2},\end{aligned} \tag{EC.9}$$

where geometric sums are obtained by assuming (without loss of generality) that $\beta_i \neq 1$.

Thus, from the perspective of sensor i, the current state S_t follows a normal distribution with mean and variance given by (EC.8) and (EC.9), respectively. Next, upon realizing the current signal $\Gamma_{it} = \gamma_{it}$, sensor i updates its prior belief about the current state according to Bayes' rule. Since both the signal received about the state and the prior on the state have a normal distribution, it follows that sensor i's posterior belief about the state is also normally distributed but with a mean and variance given by

$$\begin{aligned}\mathbb{E}[S_{it}|\Gamma_{it} = \gamma_{it}] &= \frac{\frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2n}}[\hat{\alpha}_i \frac{1-\hat{\beta}_i^n}{1-\hat{\beta}_i} + \hat{\beta}_i^n s_{t-n}] + q_i^2 \gamma_{it}}{\frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2n}} + q_i^2} \\ &= \frac{g_n(\hat{\beta}_i)[\hat{\alpha}_i \frac{1-\hat{\beta}_i^n}{1-\hat{\beta}_i} + \hat{\beta}_i^n s_{t-n}] + q_i^2 \gamma_{it}}{g_n(\hat{\beta}_i) + q_i^2}\end{aligned} \tag{EC.10}$$

and

$$
\begin{aligned}
Var[S_{it}|\Gamma_{it}=\gamma_{it}] &= \frac{1}{\frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2n}}+q_i^2} \\
&= \frac{1}{g_n(\hat{\beta}_i)+q_i^2},
\end{aligned}
\tag{EC.11}
$$

where

$$
g_n(\hat{\beta}_i) \triangleq \frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2n}}.
$$

Thus, it can be seen that the higher the quality of sensor i, the more weight it places on its signal when updating its mean belief, and the larger the associated variance reduction. In addition, assuming that the process is stationary ($|\beta_i|<1$), the smaller the n (i.e., the more recent the state observation) the smaller both the variance of the prior and the variance of posterior (see (EC.9) and (EC.11)). The special case with $n=1$ represents the case where such variances are the smallest and can be obtained by noting that $g_1(\hat{\beta}_i)=1$. Replacing $g_1(\hat{\beta}_i)=1$, yields the results obtained in the main body. However, using (EC.10) and (EC.11) with a general $g_n(\hat{\beta}_i)\triangleq\frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2n}}$ instead of their counterparts in the main body ((4) and (5), respectively) generalizes the results obtained.

Specifically, next we show that the main results obtained in the main body hold under this extension. Thus, our results reveal that our main results are *robust* to the assumption (made in the main body) that at the beginning of each period sensors have access to the realized value of the previous period's state.

Robustness of the Main Results. To study the target selection problem of sensor i, suppose it communicates with some sensor j in period t. Similar to the main body, given its privately generated signal Γ_{jt} in period t, sensor j provides sensor i with its best current estimate of state, which is $\mathbb{E}[S_{jt}|\Gamma_{jt}]$, i.e., its updated/latest expected belief about the current state S_t. Now, from sensor i's perspective, $\mathbb{E}[S_{jt}|\Gamma_{jt}]$ is formed according to:

$$
\mathbb{E}[S_{ijt}|\Gamma_{jt}] = \frac{g_n(\hat{\beta}_{ijt})[\hat{\alpha}_{ijt}\frac{1-\hat{\beta}_{ijt}^n}{1-\hat{\beta}_{ijt}}+\hat{\beta}_{ijt}^n s_{t-n}]+q_j^2\Gamma_{jt}}{g_n(\hat{\beta}_{ijt})+q_j^2}.
\tag{EC.12}
$$

Equation (EC.12) is based on (EC.10). However, it uses parameters $\hat{\alpha}_{ijt}$ and $\hat{\beta}_{ijt}$, which are sensor i's beliefs about sensor j's inference parameters $\hat{\alpha}_j$ and $\hat{\beta}_j$. (Also, note that (EC.12) reduces to (7) of the main body if we set $n=1$.) Since $\Gamma_{jt}=S_t+\epsilon_{jt}$, it can be seen from (EC.12) that $\mathbb{E}[S_{ijt}|\Gamma_{jt}]$ provides sensor i with the following noisy signal regarding the state S_t:

$$
\frac{g_n(\hat{\beta}_{ijt})+q_j^2}{q_j^2}\mathbb{E}[S_{ijt}|\Gamma_{jt}] = S_t+\epsilon_{jt}+\frac{g_n(\hat{\beta}_{ijt})[\hat{\alpha}_{ijt}\frac{1-\hat{\beta}_{ijt}^n}{1-\hat{\beta}_{ijt}}+\hat{\beta}_{ijt}^n s_{t-n}]}{q_j^2}
\tag{EC.13}
$$

There are two independent sources of noise in this signal: (a) the inherent white noise ϵ_{jt} in sensor j's measurement Γ_{jt} (which has a variance of $1/q_j^2$), and (b) the noise caused by sensor i's lack of familiarity with sensor j's inference model (the third term in the right hand side of (EC.13)). For notational convenience, we define the random variable

$$
\tilde{\Xi}_{ijt}(s_{t-n}) = g_n(\hat{\beta}_{ijt})\big[\hat{\alpha}_{ijt}\frac{1-\hat{\beta}_{ijt}^n}{1-\hat{\beta}_{ijt}}+\hat{\beta}_{ijt}^n s_{t-n}\big],
$$

and let its precision be denoted by $\tilde{\psi}_{ijt}(s_{t-n}) \triangleq 1/Var\left[\tilde{\Xi}_{ijt}(s_{t-n})\right]$. With these, based on (EC.13), the variance of the signal's noise (if i targets j in period t) is

$$\tilde{\sigma}_t^2(i,j,s_{t-n}) = \frac{q_j^2 + 1/\tilde{\psi}_{ijt}(s_{t-n})}{q_j^4},$$

which provides an extension for Equation (10) of the main body.

Similar to the main body, the target selection is then as follows. Sensor i targets the sensor that has the minimum $\tilde{\sigma}_t^2(i,j,s_{t-n})$. That is,

$$\begin{aligned} j_{it}^* &\triangleq \arg\min_{j\in\mathcal{N}\setminus\{i\}} \tilde{\sigma}_t^2(i,j,s_{t-n}) \\ &= \arg\min_{j\in\mathcal{N}\setminus\{i\}} \left\{\frac{q_j^2 + 1/\tilde{\psi}_{ijt}(s_{t-n})}{q_j^4}\right\}. \end{aligned} \tag{EC.14}$$

Now that the target selection rule is defined, we are ready to show that the main results presented in the main body hold. In particular, we show that when sensor qualities differ across sensors, for any given sensor i and along any fixed sample path $\mathcal{S}_\infty$: (a) the set of long run contacts $\mathcal{T}_i(\mathcal{S}_\infty)$ is a singleton, i.e., $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$, and (b) the unique long run contact in $\mathcal{T}_i(\mathcal{S}_\infty)$ can be identified in the almost sure sense in finite time, i.e., $\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S}_\infty^{\prime t^*})$ *a.s.* for some $t^* < \infty$. These two results, in turn, mean that one can fully define the long run communication network as a *random directed graph*, i.e., a directed graph with given probabilities assigned to each link ij that indicate the probability that j will be the long run target for i. Furthermore, there exists a finite time after which the graph can be defined as a deterministic directed graph, i.e., with all probabilities being zero or one, that fully specifies the long run target for each sensor.

To establish these results for the extension under consideration, we start by establishing the following result which extends Lemma 2 of the main body.

LEMMA EC.2. *For any $\epsilon > 0$, there exists a fixed threshold $\bar{\psi}_\epsilon \in \mathbb{R}$ such that if $\tilde{\psi}_{ijt}(s_{t-n}) > \bar{\psi}_\epsilon$ and $\frac{q_j}{q_{j'}} > 1+\epsilon$ then $j_{it}^* \neq j'$.*

Lemma (EC.2) allows us to show that the set of long run contacts of any sensor is a singleton, and that this result holds regardless of the duration of delay in state realization (i.e., parameter n). That is, Proposition 4 of the main body holds not only when $n = 1$, but also for any general n.

PROPOSITION EC.4 **(Unique Long Run Contact)**. *If sensor qualities are asymmetric, then given any sample path $\mathcal{S}_\infty$, $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$ for all $i \in \mathcal{N}$.*

Finally, we can show that another main result presented in the main body holds for the extension under consideration: when the sensor qualities are asymmetric, the long run set of contacts of each sensor can be determined in finite time. This means that transient analysis is sufficient for characterizing the communication network that will be formed in the long run, even when state realization has a general delay. Intuitively, this is because the role of state in target selection eventually vanishes: the effect of familiarity built via past communications eventually outweighs the role of state.

PROPOSITION EC.5 **(Transient Analysis)**. *If sensor qualities are asymmetric, then along any sample path $\mathcal{S}_\infty$ there exists a finite period t^* such that for all $i \in \mathcal{N}$*

$$\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S}_\infty^{\prime t^*}) \quad a.s.$$

Proposition EC.6 **(Deterministic Random Directed Graph)**. *If sensor qualities are asymmetric, then there exists a finite time t^* such that given the sample path up to t^* (i.e., $\mathcal{S}_{t^*}$), the long run communication network can be defined as a graph $\vec{G}(\mathcal{N},\mathcal{E},\mathcal{P})$ with the additional property that $p_{ij} \in \{0,1\}$ for all $i,j \in \mathcal{N}$.*

Case II: Delayed-and-Imperfect State Observations

We next consider the case in which each sensor's slower-but-more-accurate estimation approach has a general delay of n periods but is imperfect. In other words, the delayed state observations are subject to errors. To extend Case I, we assume that right before period t starts, the value of state at $t-n$ (for some positive integer $n \le t$) is revealed with some random error that is independent across sensors. Specifically, we model this revelation for sensor i as $S'_{i,t-n} = S_{t-n} + \upsilon_i$ with a realized value denoted by $S'_{i,t-n} = s'_{i,t-n}$, where S_{t-n} is the true value of state at period $t-n$, and υ_i represents a normally distributed random noise with mean 0 and variance $1/q_{i0}^2$. In this setting, q_{i0}^2 denotes the precision (i.e., quality) of sensor i's slower-but-more-accurate approach in measuring the state. For expositional ease, we will assume this quality is the same across sensors, that is, $q_{i0} = q_0\ \forall i$. Because the slower approach is more accurate than the regular sensors, it follows that $q_0 >> q_i\ \forall i$. Of note, if $q_0 = \infty$ this setting is equivalent to Case I studied earlier. However, we now generalize our previous analyses by allowing $q_0 < \infty$. We assume the υ_i noise terms are independent across sensors. However, the case in which there is one "universal" slower-but-more accurate approach that gives the same (imperfect) measurement to all sensors is easily handled by simply setting $\upsilon_i = \upsilon\ \forall i$ in what follows.

Recall from (EC.7) that

$$S_{it} = \hat{\alpha}_i \sum_{k=0}^{n-1} \hat{\beta}_i^k + \hat{\beta}_i^n S_{t-n} + \sum_{k=0}^{n-1} \hat{\beta}_i^k \tilde{\epsilon}_{t-k}. \tag{EC.15}$$

Hence, we have

$$\begin{aligned}\mathbb{E}[S_{it}|S'_{i,t-n} = s'_{i,t-n}, S'_{i,t-n} = S_{t-n} + \upsilon_i] &= \hat{\alpha}_i \sum_{k=0}^{n-1} \hat{\beta}_i^k + \hat{\beta}_i^n s'_{i,t-n} \\ &= \hat{\alpha}_i \frac{1-\hat{\beta}_i^n}{1-\hat{\beta}_i} + \hat{\beta}_i^n s'_{i,t-n}\end{aligned} \tag{EC.16}$$

and

$$\begin{aligned}Var[S_{it}|S'_{i,t-n} = s'_{i,t-n}, S'_{i,t-n} = S_{t-n} + \upsilon_i] &= \frac{\hat{\beta}_i^{2n}}{q_0^2} + \sum_{k=0}^{n-1} \hat{\beta}_i^{2k} Var[\tilde{\epsilon}_{t-k}] \\ &= \frac{\hat{\beta}_i^{2n}}{q_0^2} + \frac{1-\hat{\beta}_i^{2n}}{1-\hat{\beta}_i^2},\end{aligned} \tag{EC.17}$$

where geometric sums are obtained by assuming (without loss of generality) that $\beta_i \neq 1$.

Thus, similar to Case I, from the perspective of sensor i, the current state S_t follows a normal distribution with mean and variance given by (EC.16) and (EC.17), respectively. Next, upon realizing the current signal $\Gamma_{it} = \gamma_{it}$, sensor i updates its prior belief about the current state according to Bayes' rule. Since both the signal received about the state and the prior on the state have a normal distribution, it follows that sensor i's posterior belief about the state is also normally distributed but with a mean and variance given by

$$\mathbb{E}[S_{it}|\Gamma_{it}=\gamma_{it}] = \frac{\big[\frac{\hat{\beta}_i^{2n}}{q_0^2}+\frac{1-\hat{\beta}_i^{2n}}{1-\hat{\beta}_i^2}\big]^{-1}[\hat{\alpha}_i\frac{1-\hat{\beta}_i^n}{1-\hat{\beta}_i}+\hat{\beta}_i^n s'_{i,t-n}]+q_i^2\gamma_{it}}{\big[\frac{\hat{\beta}_i^{2n}}{q_0^2}+\frac{1-\hat{\beta}_i^{2n}}{1-\hat{\beta}_i^2}\big]^{-1}+q_i^2}$$
$$= \frac{g_n(\hat{\beta}_i)[\hat{\alpha}_i\frac{1-\hat{\beta}_i^n}{1-\hat{\beta}_i}+\hat{\beta}_i^n s'_{i,t-n}]+q_i^2\gamma_{it}}{g_n(\hat{\beta}_i)+q_i^2} \tag{EC.18}$$

and

$$Var[S_{it}|\Gamma_{it}=\gamma_{it}] = \frac{1}{\big[\frac{\hat{\beta}_i^{2n}}{q_0^2}+\frac{1-\hat{\beta}_i^{2n}}{1-\hat{\beta}_i^2}\big]^{-1}+q_i^2}$$
$$= \frac{1}{g_n(\hat{\beta}_i)+q_i^2}, \tag{EC.19}$$

where

$$g_n(\hat{\beta}_i) \triangleq \big[\frac{\hat{\beta}_i^{2n}}{q_0^2}+\frac{1-\hat{\beta}_i^{2n}}{1-\hat{\beta}_i^2}\big]^{-1}.$$

When $q_0=\infty$ we observe that $g_n(\hat{\beta}_i)$ given above is equivalent to that obtained under Case I. However, it provides a generalization by allowing $q_0<\infty$. Importantly, using the $g_n(\hat{\beta}_i)$ given above instead of that in Case I, it can easily been shown that all the results provided under Case I still hold. This follows directly if the slower-but-more accurate approach is common to all sensors (i.e., they all receive the same delayed-but-imperfect measurement). If the delayed-but-imperfect measurements are independent then the results also go through directly subject to a mild assumption that sensors do not try to learn about the delayed measurements received by other sensors. Any benefits of such learning (to improve estimation) would be very low (and so can be ignored) because the slower-but-more-accurate estimation approach has a much higher precision than the real-time sensor measurement, i.e., $q_0 >> q_i$.

Case III: No State Observations, Transient Behavior

In the previous subsection, we considered the case where state is realized with some delay (Case I) and showed that the main results presented in the main body hold under such an extension. We now consider the case in which there is no state observation beyond the one in the very first period, s_0 (see below for another extension in which even s_0 is not known). This case can be analyzed by setting $n=t$ and using the result of Case I. Specifically, for $n=t$, we observe that

$$\mathbb{E}[S_{it}|\Gamma_{it}=\gamma_{it}] = \frac{\frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2t}}[\hat{\alpha}_i\frac{1-\hat{\beta}_i^t}{1-\hat{\beta}_i}+\hat{\beta}_i^t s_0]+q_i^2\gamma_{it}}{\frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2t}}+q_i^2}$$
$$= \frac{g_t(\hat{\beta}_i)[\hat{\alpha}_i\frac{1-\hat{\beta}_i^t}{1-\hat{\beta}_i}+\hat{\beta}_i^t s_0]+q_i^2\gamma_{it}}{g_t(\hat{\beta}_i)+q_i^2} \tag{EC.20}$$

and

$$Var[S_{it}|\Gamma_{it}=\gamma_{it}] = \frac{1}{\frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2t}}+q_i^2}$$
$$= \frac{1}{g_t(\hat{\beta}_i)+q_i^2}, \tag{EC.21}$$

where

$$g_t(\hat{\beta}_i) \triangleq \frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2t}}.$$

Similar to Case I, using (EC.20) and (EC.21) with a general $g_t(\hat{\beta}_i) \triangleq \frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2t}}$ instead of their counterparts in the main body ((4) and (5), respectively) generalizes the results obtained (after some necessary modifications).

What if even s_0 is not available? Note that the availability of s_0 in the above analyses is not restrictive. For example, suppose s_0 is not available to the sensors but they either make use of the historical mean of the state value (e.g., using their training observations) to find/estimate a proxy for s_0, or rely on a predicated value provided by an outside entity (e.g., a third-party's smart device). In either case, the s_0 value in the above analysis can be replaces with $\tilde{s}_0$, where $\tilde{s}_0$ represents the proxy available to the sensors. Replacing s_0 with its proxy $\tilde{s}_0$ (which might differ from s_0), we have

$$\begin{aligned}\mathbb{E}[S_{it}|\Gamma_{it}=\gamma_{it}] &= \frac{\frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2t}}[\hat{\alpha}_i\frac{1-\hat{\beta}_i^t}{1-\hat{\beta}_i}+\hat{\beta}_i^t\tilde{s}_0]+q_i^2\gamma_{it}}{\frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2t}}+q_i^2} \\ &= \frac{g_t(\hat{\beta}_i)[\hat{\alpha}_i\frac{1-\hat{\beta}_i^t}{1-\hat{\beta}_i}+\hat{\beta}_i^t\tilde{s}_0]+q_i^2\gamma_{it}}{g_t(\hat{\beta}_i)+q_i^2} \end{aligned} \tag{EC.22}$$

Using using (EC.22) and (EC.21) with a general $g_t(\hat{\beta}_i) \triangleq \frac{1-\hat{\beta}_i^2}{1-\hat{\beta}_i^{2t}}$ instead of their counterparts in the main body ((4) and (5), respectively), we again see that the main results are generalizable. Furthermore, in the next case (Case III), we analyze the setting in which the system dynamics has reached steady-state. We note that in that case, even a proxy for s_0 is not needed.

Case IV: No State Observations, Steady-State Behavior

We now consider a similar case to that of Case II by assuming that there is no state observation. However, unlike Case II, we assume that at the beginning of the horizon the state dynamics has reached steady-state. This case can be analyzed by setting $n \to \infty$, and nothing that the process is stationary if $|\beta_i| < 1$. To this end, we define

$$g_\infty(\hat{\beta}_i) \triangleq 1-\hat{\beta}_i^2.$$

It can then be seen that

$$\begin{aligned}\mathbb{E}[S_{it}|\Gamma_{it}=\gamma_{it}] &= \frac{g_\infty(\hat{\beta}_i)[\hat{\alpha}_i\frac{1}{1-\hat{\beta}_i}]+q_i^2\gamma_{it}}{g_\infty(\hat{\beta}_i)+q_i^2} \\ &= \frac{(1+\hat{\beta}_i)\hat{\alpha}_i+q_i^2\gamma_{it}}{g_\infty(\hat{\beta}_i)+q_i^2}\end{aligned} \tag{EC.23}$$

and

$$Var[S_{it}|\Gamma_{it}=\gamma_{it}] = \frac{1}{g_\infty(\hat{\beta}_i)+q_i^2}. \tag{EC.24}$$

Again, similar to cases I and II, using (EC.23) and (EC.24) with $g_\infty(\hat{\beta}_i) \triangleq 1-\hat{\beta}_i^2$ instead of their counterparts in the main body ((4) and (5), respectively) generalizes the results obtained (after some necessary modifications).

Appendix D: Correlated Measurement Errors

In the main body, we assumed that sensors take independent readings of the state variable. In particular, we assumed that each sensor i receives a noisy signal about the state variable at in time t:

$$\Gamma_{it} = S_t + \epsilon_{it}, \tag{EC.25}$$

where ϵ_{it} are i.i.d. normal white noises with mean 0 and variance $1/(q_i)^2$, with q_i representing the quality of sensor i. In this setting, a higher quality sensor has a higher precision. Furthermore, it should be noted that even in this setting sensor measurements given by (EC.25) are highly correlated. Indeed, in high states, all sensor readings tend to be high, and in low states, all sensor readings are low. Moreover, since the value of state is correlated across periods, sensor readings are correlated even across periods.

However, in the main body, we assumed that the errors in these readings are not correlated across sensors. That is, we assumed that ϵ_{it} and ϵ_{jt} are independent for any $i \neq j$. Here, we extend our results by allowing these errors to be correlated.

In what follows, we discuss two cases separately. The first case is where the correlation in these measurement errors are the same among all sensors. We label this case as "symmetric correlated measurement errors" and show that almost all our major results in the main body remain unchanged for this case, because the criterion each sensor will use to choose its target is essentially *equivalent* to the one established in the main body. In other words, each sensor will choose the same target as it would if there was no correlation in measurement errors.

The second case, which we label as "asymmetric correlated measurement errors," will consider the situation in which the amount of correlation in measurement errors between a sensor i and a sensor j depends on i and j (thus, asymmetric across sensors). We show that this case induces a target selection criterion that is in general *not equivalent* to the one established in the main body. Hence, in general, a sensor i might communicate to a different sensor than it would if there was no correlation. However, we are able to establish conditions under which even in this case the target selection criterion of each sensor will be equivalent to that provided in the main body. That is, under some conditions, we demonstrate that, even in this case, each sensor will choose in each period exactly the same sensor that it would when the errors in measurements are not correlated.

Notably, both of the above-mentioned cases shed light on the importance of understanding the setting in which the measurement errors are uncorrelated, which is the focus of the main body of the paper.

Symmetric Correlated Measurement Errors. Consider the case in which the error terms in sensor readings given by (EC.25) are correlated in a symmetric way. Specifically, for any two sensors i and j with $i \neq j$ assume that

$$\epsilon_{jt} = \nu\, \epsilon_{it} + \epsilon'_{jt}, \tag{EC.26}$$

where ϵ'_{jt} are i.i.d. normal white noises with mean 0 and variance $1/(q_i)^2$. In this setting, the level of correlation in measurement errors is denoted by ν, which does not depend on i and j. Similar to (8) in the main body, we have that if sensor i communicates with sensor j then it receives the following noisy signal regarding the state S_t:

$$\frac{1+q_j^2}{q_j^2}\mathbb{E}[S_{ijt}|\Gamma_{jt}] = S_t + \epsilon_{jt} + \frac{\hat{\alpha}_{ijt} + \hat{\beta}_{ijt} s_{t-1}}{q_j^2}. \tag{EC.27}$$

Similar to (9), we can denote the variance in this signal's noise as:

$$\sigma_t'^2(i,j,s_{t-1}) = Var\big[\epsilon_{jt} + \frac{\hat{\alpha}_{ijt} + \hat{\beta}_{ijt}s_{t-1}}{q_j^2}\big]. \tag{EC.28}$$

Replacing (EC.26) in (EC.30), we have

$$\sigma_t'^2(i,j,s_{t-1}) = \frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4} + \frac{\nu^2}{q_i^2}, \tag{EC.29}$$

where similar to the main body $\psi_{ijt}(s_{t-1}) \triangleq 1/Var\big[\Xi_{ijt}(s_{t-1})\big]$. Comparing (EC.29) with (10), we observe that

$$\sigma_t'^2(i,j,s_{t-1}) = \sigma_t^2(i,j,s_{t-1}) + \frac{\nu^2}{q_i^2}, \tag{EC.30}$$

where $\sigma_t^2(i,j,s_{t-1})$ in the variance of the signal i will receive from j (if they communicate at time t). From (EC.30), we can establish the following equivalence result, which shows that each sensor will choose the same target whether $\nu = 0$ or $\nu \neq 0$. This, in turn, implies that the main results developed in the main body hold, and the network that will be formed among sensors will be the same as the one studied there (both in short-term and in long-term).

PROPOSITION EC.7. *Let j_{it}^* and $j_{it}'^*$ denote the optimal target of sensor i in period t when $\nu = 0$ and $\nu \neq 0$, respectively (all else equal). Then $j_{it}^* = j_{it}'^*$.*

Since the above proposition establishes that every sensors in each period would select the same sensor whether $\nu = 0$ or $\nu \neq 0$, when all else (including the sample path of state realizations) is equal, all the main results of the main body still hold for this scenario.

Asymmetric Correlated Measurement Errors. Now consider the scenario in which correlations in measurement errors are not symmetric. Specifically, for any two sensors i and j with $i \neq j$ assume that

$$\epsilon_{jt} = \nu_{ij}\,\epsilon_{it} + \epsilon'_{jt}, \tag{EC.31}$$

where ϵ'_{jt} are i.i.d. normal white noises with mean 0 and variance $1/(q_i)^2$. In this setting, the level of correlation in measurement errors is denoted by ν_{ij}, which depends on i and j. Similar to the previous case, we can denote the variance in this signal's noise received by i if it communicates to j in period t by:

$$\sigma_t'^2(i,j,s_{t-1}) = Var\big[\epsilon_{jt} + \frac{\hat{\alpha}_{ijt} + \hat{\beta}_{ijt}s_{t-1}}{q_j^2}\big]. \tag{EC.32}$$

Replacing (EC.31) in (EC.32), we have

$$\sigma_t'^2(i,j,s_{t-1}) = \frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4} + \frac{\nu_{ij}^2}{q_i^2}, \tag{EC.33}$$

where similar to the main body $\psi_{ijt}(s_{t-1}) \triangleq 1/Var\big[\Xi_{ijt}(s_{t-1})\big]$. Comparing (EC.33) with (10), we observe that

$$\sigma_t'^2(i,j,s_{t-1}) = \sigma_t^2(i,j,s_{t-1}) + \frac{\nu_{ij}^2}{q_i^2}, \tag{EC.34}$$

where $\sigma_t^2(i, j, s_{t-1})$ is the variance of the signal i will receive from j (if it communicate to j time t) when the measurement errors are independent. Since the second term in (EC.34) depends on j, we immediately observe that target selection in general in this case is not equivalent to the one studied in the main body. However, as we show below, we are able to establish conditions under which a similar result to the equivalence result in Proposition EC.7 holds.

In particular, the following proposition shows that if the ratio of the maximum of the differences in correlation values across sensors to the quality of sensor i is lower than a threshold, then sensor i will choose the same target whether $\nu_{ij} = 0$ or $\nu_{ij} \neq 0$. Since correlation values typically change in a limited range (e.g., [-1,1]), and quality of a sensor i given by q_i is typically relatively large, this sufficient condition established in the result below is not very strong.

PROPOSITION EC.8. *Let j_{it}^* and $j_{it}'^*$ denote the optimal target of sensor i in period t when $\nu_{ij} = 0$ and $\nu_{ij} \neq 0$, respectively (all else equal). There exists a threshold $0 < k_t < \infty$ such that if*

$$\max_{j,j' \neq i} \frac{\nu_{ij}^2 - \nu_{ij'}^2}{q_i^2} < k_t,$$

then $j_{it}^* = j_{it}'^*$.

Note:
Appendix E (Non-Myopic Behavior), Appendix F (Self-Updating Inference Models), Appendix G (Proofs of Results from Main Paper), Appendix H (Proofs of Results from Appendices), and Appendix I (Numerical Studies Detailed Description and Discussion) are not part of the basic e-companion for reasons of space. They are available in the unabridged version of the paper posted on SSRN.

Appendix E: Non-Myopic Behavior

In the main body, we assumed that sensors were myopic: in each period, each sensor i took the action that minimized its estimation squared error of that period:

$$\begin{aligned} \min_{\tilde{s}_{it}\in\mathbb{R},\mathbf{a}_{it}\in\{0,1\}^{n-1}} \quad & \mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}\left[\tilde{s}_{it}-S_t\right]^2 \\ s.t. \quad & \\ & 0<c\textstyle\sum_{j\in\mathcal{N}\setminus\{i\}} a_{ijt}\leq b, \end{aligned} \tag{EC.35}$$

where $\lfloor\frac{b}{c}\rfloor=1$, and the vector $\mathbf{a}_{it}\in\{0,1\}^{n-1}$ is composed of elements a_{ijt} with $a_{ijt}=1$ if i targets j, and $a_{ijt}=0$ otherwise. In the proof of Proposition 1, we showed that this is equivalent to minimizing the posterior variance $Var_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t]$, where if $a_{ijt}=1$, the posterior variance is

$$Var_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t]=\frac{1}{1+q_i^2+[\sigma_t^2(i,j,s_{t-1})]^{-1}}\,, \tag{EC.36}$$

with

$$\sigma_t^2(i,j,s_{t-1}) = \frac{q_j^2+1/\psi_{ijt}(s_{t-1})}{q_j^4}. \tag{EC.37}$$

We now consider an extension of this setting in which sensors are not myopic. When a sensor is not myopic, it might be willing to communicate with a sensor that does not provide it with the best information regarding the current state if doing so helps it build familiarity with that sensor for future periods.

To analyze the extension, based on (EC.36), we first define the following disutility function for sensor i:

$$\phi_i(\mathbf{a}_{it})\triangleq Var_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t]=\sum_{j\in\mathcal{N}\setminus\{i\}}\frac{a_{ijt}}{1+q_i^2+[\sigma_t^2(i,j,s_{t-1})]^{-1}}. \tag{EC.38}$$

We then assume that each sensor i minimizes the discounted sum of its disutilities over the entire horizon $t=1,2,\cdots,T$:

$$\begin{aligned} \min_{\mathbf{a}_{it}\in\{0,1\}^{n-1},\forall t=1,\cdots,T} \quad & \mathbb{E}\Big[\textstyle\sum_{t=1}^{T}\eta^{t-1}\,\phi_i(\mathbf{a}_{it})\Big] \\ s.t. \quad & \\ & \textstyle\sum_{j\in\mathcal{N}\setminus\{i\}} a_{ijt}=1, \qquad \forall t=1,\cdots,T, \end{aligned} \tag{EC.39}$$

where $\eta\in[0,1)$ is a discount factor, the constraint ensures that exactly one sensor can be targeted at each period, and the expectation operator in (EC.39) is with respect to the sequential revelation of states. Hence, each sensor minimizes the discounted sum of its posterior variance, which is equivalent to minimizing the discounted sum of its estimation squared errors (i.e., the discounted sum of the values of its loss function). The problem in infinite-horizon can be analyzed by considering (EC.39) and taking the limit as $T\to\infty$.

To gain some insight, let $\mathbf{m}=(m_j)_{j\in\mathcal{N}\setminus\{i\}}$ denote sensor i's vector of the number of previous communications with other sensors and let s denote the last state realization. Then, the optimization program (EC.39) can be solved via the following dynamic program:

$$V_t^i(\mathbf{m},s)=\min_{j\in\mathcal{N}\setminus\{i\}}\left\{\phi_{ij}(m_j,s)+\eta\int V_{t-1}^i(\mathbf{m}+\mathbf{e}_j,s')\,dF_{s'|s}^i(s')\right\} \tag{EC.40}$$

along with the terminal condition $V_0^i\big(\mathbf{m}, s\big) = 0$, where $V_t^i(\mathbf{m}, s)$ represents the value function (disutility) if there are t periods to go. In (EC.40), the immediate "cost" based on (EC.36) and (EC.37) is

$$\phi_{ij}(m_j, s) \triangleq \frac{1}{1 + q_i^2 + \frac{q_j^4}{q_j^2 + 1/\psi_{ij}(s)}}, \tag{EC.41}$$

where based on Proposition 2

$$\psi_{ij}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j) m_j\,. \tag{EC.42}$$

Furthermore, in (EC.40), $\mathbf{e}_j$ denotes a vector with the element associated with sensor j equal to one and all other elements equal to zero, and $F^i_{s'|s}(s')$ represents sensor i's belief about the next state (i.e., s') given that the last state realization is s.

The information state in (EC.40) is the number of previous communications along with the last revealed state. If sensor i chooses to connect to sensor j, it incurs the immediate "cost" of $\phi_{ij}(m_j, s)$, its number of communications with sensor j increases by one (equivalently, its familiarity to sensor j improves by $f(q_i, q_j)$; see Proposition 2), the revealed state of the environment next changes from s to s', and there remains $t-1$ periods to go.

Although (EC.40) is presented for a finite horizon setting ($T < \infty$), taking limit as $t \to \infty$ allows studying the problem in an infinite-horizon scenario ($T = \infty$). We denote the infinite-horizon value function for sensor i by $V_\infty^i\big(\mathbf{m}, s\big)$, and based on $(EC.41)$ and $(EC.42)$ let

$$\phi_{ij}(\infty, s) \triangleq \lim_{m_j \to \infty} \phi_{ij}(m_j, s) = \frac{1}{1 + q_i^2 + q_j^2}.$$

Of note, (EC.40) can be viewed as a *restless multi-armed bandit* problem, if we think of each sensor as an arm. Each arm j has an underlying Markov chain with state (m_j, s). The transitions in this Markov chain depends on whether the arm is pulled or not. Even if an arm is not pulled, its state changes, and hence, the bandits are *restless*.

In general, *restless multi-armed bandit* problems are hard to analyze. Particularly, (a) it is known that a problem in this class might not be *indexable*, and (b) there no known general conditions that guarantee optimality of a *myopic* policy. In what follows, we seek to generate insights into conditions under which each sensor i will act myopically (i.e., by connecting to the sensor that has the minimum immediate "cost" $\phi_{ij}(m_j, s)$). Under such conditions, all the results established in the main body will hold for non-myopic sensors as well.

Optimality of Myopic Behavior We now provide some structural properties and establish some conditions under which its optimal for sensors to act myopically. We start by noting that since the immediate "cost" function $\phi_{ij}(m_j, s)$ is decreasing in m_j for all s, we have

$$\lim_{m_j \to \infty} \phi_{ij}(m_j, s) \le \phi_{ij}(m_j, s) \le \lim_{m_j \to 0} \phi_{ij}(m_j, s). \tag{EC.43}$$

Replacing the limits in above, we observe that

$$\frac{1}{1 + q_i^2 + q_j^2} \le \phi_{ij}(m_j, s) \le \frac{1}{1 + q_i^2 + \frac{q_j^4}{q_j^2 + \frac{w_{ij0}^2 + v_{ij0}^2 s^2}{v_{ij0}^2 w_{ij0}^2}}}. \tag{EC.44}$$

Hence, the immediate "cost" function $\phi_{ij}(m_j, s)$ is bounded. This along with the fact that there are finite number of sensors allows us to establish the following intuitive result.

Proposition EC.9 **(Threshold on η)**. *There exists a threshold $\bar{\eta}(\mathbf{m}, s) > 0$ such that if $\eta \leq \bar{\eta}(\mathbf{m}, s)$, then the optimal policy for all sensors at information state $(\mathbf{m}, s)$ in both finite-horizon and infinite-horizon settings is myopic.*

Proposition EC.9 states the following intuitive condition: if the discount factor is below a certain threshold, then it is optimal for sensors to act myopically. This might happen when the length of each period is relatively long, or when good current estimates are much more valuable than those in future. To dig deeper, we next examine conditions under which for any given discount factor (i.e., even when the discount factor is not relatively low) it is optimal for sensors to act myopically.

In particular, we next demonstrate that there exists a state-dependent threshold on m_j denoted by $\bar{m}_j$ (after suppressing its dependency to other elements of information state for the ease of notation) such that, if it is optimal for sensor i to choose sensor j myopically when its number of previous communication with j is $\bar{m}_j$, then it is optimal for sensor i to do the same whenever $m_j \geq \bar{m}_j$ (keeping all else the same).

Proposition EC.10 **(Threshold on m_j)**. *Suppose at some information state $(\mathbf{m}^*, s)$ it is optimal for sensor i to myopically choose to communicate to sensor j. Then it is optimal for sensor i to do the same at $(\mathbf{m}^* + \mathbf{e}_j, s)$ in both finite-and infinite-horizon settings. Furthermore, there exists a threshold $\bar{m}_j$ such that it is optimal for sensor i to myopically choose to communicate to sensor j at any information state $(\mathbf{m}, s)$ with $m_j \geq \bar{m}_j$ and $m_{j'} = m^*_{j'}$ for all $j' \in \mathcal{N} \setminus \{i, j\}$.*

Propositions EC.9 and EC.10 establish sufficient conditions under which it is optimal for sensors to act myopically. However, these conditions are not necessary, and hence, there might be various other situations in which it is optimal for sensors to choose their targets myopically. It should be noted, however, that a myopic behavior is not always optimal. Nevertheless, the insights gained in the main body by focusing on the case where sensor communications are myopic are highly informative as they shed light on how networks may form among sensors as well as on the related tradeoffs that may arise in IoT.

Appendix F: Self-Updating Inference Models

In the main body, we assumed that each sensor i had its own *unchanging* inference model about the environment. In particular, while the true state dynamics (given in (1)) is

$$S_t = \alpha + \beta S_{t-1} + \tilde{\epsilon}_t \tag{EC.45}$$

we assumed that sensor i's model is

$$S_{it} = \hat{\alpha}_i + \hat{\beta}_i S_{t-1} + \tilde{\epsilon}_t, \tag{EC.46}$$

where $\hat{\alpha}_i$ and $\hat{\beta}_i$ are sensor i's estimates of the true process parameters α and β, and $\tilde{\epsilon}_t$ is a normal white-noise random variables with mean 0 and variance normalized to 1.

Firms typically use their own training data sets and training algorithms prior to deployment, and hence, sensors developed by different firms will have different inference models. Because firms typically use massive training data sets, we assumed in the main body that once the training phase ends and deployment begins, firms do not adjust their inference models. This is reasonable when the amount of the data used during training is much larger than the data that can be collected during deployment, and thus, the effect of using each single data point collected during deployment to dynamically update the estimates of the α and β parameters is negligible. Of note, even with a large training data set, each sensor's estimate of environment parameters α and β is rarely perfect, especially in situations in which it is important to use a good training algorithm. For such reasons, estimates of α and β are typically (a) firm-dependent, and (b) imperfect. Nevertheless, since such estimates are based on much larger data sets than those that can be obtained during the deployment process, it is reasonable for firms to ignore minor updating of their estimates during the deployment phase, which is the setting we studied in the main body.

However, below we extend our analysis by allowing each sensor to update its estimates after each state realization. Recall that, in our main setting, right before each period starts the state of the previous period is realized (see Online Appendix B for extensions in which there is no state realization or there are delayed state realizations). To see the impact of using this state realization to update estimates of α and β parameters, suppose that prior to the realization of s_t, sensor i believes

$$(\alpha_{it}, \beta_{it})^T \sim_i N\big((\alpha, \beta)^T + \mathbf{\Delta}_{it}, \mathbf{\Lambda}_{it}^{-1}\big).$$

That is, his parameters $(\alpha_{it}, \beta_{it})^T$ (considered as a column vector) are random variables having a bivariate normal distribution with a mean estimation error given by the (column) vector $\mathbf{\Delta}_{it}$ and precision $\mathbf{\Lambda}_{it}$. Sensor i's estimated parameters prior to realization of s_t are then $(\hat{\alpha}_{it}, \hat{\beta}_{it})^T = (\alpha, \beta)^T + \mathbf{\Delta}_{it}$. When s_t is realized, sensor i makes the following observation about $(\alpha_{it}, \beta_{it})^T$:

$$\alpha_{it} + \beta_{it} s_{t-1} + \tilde{\epsilon}_t = s_t.$$

Assuming that sensor i uses Bayesian updating, it follows from the general result on multi-variate updating of normal distribution (discussed under Method 2 of Proof of Proposition 2) that its posterior over the parameters denoted by $(\alpha_{i,t+1}, \beta_{i,t+1})^T$ becomes

$$(\alpha_{i,t+1}, \beta_{i,t+1})^T \sim_i N(\mathbf{\Sigma}\mathbf{A}^T s_t + \mathbf{\Sigma}\mathbf{\Lambda}_{it}\big[(\alpha, \beta)^T + \mathbf{\Delta}_{it}\big], \mathbf{\Sigma}),$$

where

$$\mathbf{\Sigma}^{-1} = \mathbf{\Lambda}_{it} + \mathbf{A}^T\mathbf{A} \tag{EC.47}$$

and $\mathbf{A} = (1, s_{t-1})$. If we assume parameter estimations are based on mean of distributions (e.g., assuming that they are based on least squared errors or maximum likelihood estimation) then the change (denoted by vector $\mathbf{C}$) in sensor i's estimation of $(\alpha, \beta)^T$ due to the updating is

$$\mathbf{C} \triangleq \mathbf{\Sigma} \mathbf{A}^T s_t + \big[\mathbf{\Sigma} \mathbf{\Lambda}_{it} - \mathbf{I}\big]\big[(\alpha, \beta)^T + \mathbf{\Delta}_{it}\big] \tag{EC.48}$$

where $\mathbf{I}$ represents the identity matrix.

In what follows, we first use Lemma EC.4 (see Appendix G) to show that the change $\mathbf{C}$ can be arbitrarily small (for any state realization) when $\mathbf{\Lambda}$ is large enough. As noted earlier, $\mathbf{\Lambda}$ is typically large since firms use large amount of training data prior to deployment. We then use this result along with the fact that there are finite number of sensors to show that there exist a threshold such that if the precision of all sensors are not below it, parameter updating will not alter the target selection of any of the sensors: target selection is robust to the assumption that sensors do not update their inference parameters. Importantly, this implies that network formation among sensors can be studied without assuming that sensors update their inference models during the deployment phase.

Lemma EC.3. *For any $\epsilon > 0$ and $s = s_{t-1}$, there exists a threshold $\bar{\mathbf{\Lambda}}_{\epsilon,s}$ such that if at time t $\mathbf{\Lambda} \geq \bar{\mathbf{\Lambda}}_{\epsilon,s}$ for all sensors, then $-\epsilon \mathbf{I} \leq \mathbf{C} \leq \epsilon \mathbf{I}$ for all sensors.*

Using Lemma EC.3, we next show that if sensors have enough precisions in their inference model's parameters (e.g., have trained their sensor's using enough training data) and this fact is publicly known, then target selection is not altered by whether they update their parameters or not during the deployment phase. Hence, for the purpose of understanding how sensors choose their targets, and how communication networks form among them in the era of Big Data, it is safe to assume that there is no updating of the inference model parameters during the deployment phase.

Proposition EC.11. *Let j^*_{it} and j^{**}_{it} denote the optimal target of sensor i at time t when sensors are not allowed and are allowed to update their inference model's parameters, respectively. For any $s = s_{t-1}$, there exists a threshold $\bar{\mathbf{\Lambda}}_s$ such that if (a) $\mathbf{\Lambda} \geq \bar{\mathbf{\Lambda}}_s$ for all sensors, and (b) condition (a) is known by all sensors, then $j^*_{it} = j^{**}_{it}$ for all $i \in \mathcal{N}$.*

Appendix G: Proofs of Results from Main Paper

Proof of Proposition 1: Consider the following optimization program:

$$\min_{\tilde{s}_{it}\in\mathbb{R},\mathbf{a}_{it}\in\{0,1\}^{n-1}} \quad \mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}\left[\tilde{s}_{it}-S_t\right]^2 \qquad \text{(EC.49)}$$
$$s.t.$$
$$0 < c\sum_{j\in\mathcal{N}\setminus\{i\}} a_{ijt} \le b,$$

where $\lfloor \frac{b}{c} \rfloor = 1$, and the vector $\mathbf{a}_{it}\in\{0,1\}^{n-1}$ is composed of elements a_{ijt} with $a_{ijt}=1$ if i targets j, and $a_{ijt}=0$ otherwise. Note that each feasible solution to this optimization program has exactly one a_{ijt} equal to one. Furthermore, under each such feasible solution, it follows from the convexity of the objective function in decision variable $\tilde{s}_{it}$ that the optimal estimation for $\tilde{s}_{it}$ satisfies the first order condition

$$\frac{\partial}{\partial\tilde{s}_{it}}\left[\mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}\left[\tilde{s}_{it}-S_t\right]^2\right]=0. \qquad \text{(EC.50)}$$

From (EC.50), we observe that:

$$\tilde{s}^*_{it}=\mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t].$$

Therefore, in (EC.49), we have

$$\begin{aligned}\min_{\tilde{s}_{it}\in\mathbb{R},\mathbf{a}_{it}\in\{0,1\}^{n-1}} \mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}\left[\tilde{s}_{it}-S_t\right]^2 &= \min_{\mathbf{a}_{it}\in\{0,1\}^{n-1}} \mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}\left[\mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t]-S_t\right]^2\\ &= \min_{\mathbf{a}_{it}\in\{0,1\}^{n-1}} Var_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t], \qquad \text{(EC.51)}\end{aligned}$$

where the last term denotes sensor i's variance of S_t under the posterior distribution $F_{it}^{\mathbf{a}_{it}}$. That is, sensor i chooses its target based on minimizing the posterior variance. Next, note that j^*_{it} given by

$$\begin{aligned}j^*_{it} &\triangleq \arg\min_{j\in\mathcal{N}\setminus\{i\}} \sigma^2_t(i,j,s_{t-1})\\ &= \arg\min_{j\in\mathcal{N}\setminus\{i\}} \left\{\frac{q_j^2+1/\psi_{ijt}(s_{t-1})}{q_j^4}\right\}. \qquad \text{(EC.52)}\end{aligned}$$

minimizes the posterior variance. This is because, with $a_{ijt}=1$, the posterior variance is

$$Var_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t]=\frac{1}{1+q_i^2+[\sigma^2_t(i,j,s_{t-1})]^{-1}}, \qquad \text{(EC.53)}$$

in which $1+q_i^2$ is sensor i's precision prior to communication with sensor j, and $[\sigma^2_t(i,j,s_{t-1})]^{-1}$ is the precision of the signal it receives from sensor j. Hence, the optimal solution satisfies $a^*_{ijt}=\mathbb{1}_{\{j=j^*_{it}\}}$, and the proof is complete. □

Proof of Proposition 2: For completeness, we provide two methods of proving the result. In the first method, we only consider the aggregated univariate random variable $\Xi_{ijt}(s_{t-1})=\hat{\alpha}_{ijt}+\hat{\beta}_{ijt}s_{t-1}$ without explicitly utilizing and updating the joint distribution of α_{ijt} and β_{ijt}. In the second proof, we explicitly consider the joint distribution of α_{ijt} and β_{ijt} and update it over time as communications are made.

Method 1: First consider the following general setting. Let $Y=X+\epsilon$ be an observation made about some random variable X. It follows from Bayesian updating that if X and ϵ are both normally distributed, so is the posterior distribution for the random variable $X|Y=y$. Furthermore, it is easy to show that the precision of the random variable $X|Y=y$ is the sum of precisions of X and

ϵ. Next, let $X = \Xi_{ijt}(\cdot)$ and $Y = \Xi_{ijt}(\cdot) + q_j^2(S_t + \epsilon_{jt})$, and note that when sensor i communicates with sensor j at time t (i.e., when $a_{ijt} = 1$) the noise term in the observation (Y) made about $X = \Xi_{ijt}(\cdot)$ has variance

$$\begin{aligned} Var[q_j^2(S_t + \epsilon_{jt})] &= q_j^4\, Var[(S_t + \epsilon_{jt})] \\ &= q_j^4\,\big(\frac{1}{1+q_i^2} + \frac{1}{q_j^2}\big) \\ &= q_j^4/(1+q_i^2) + q_j^2, \end{aligned} \tag{EC.54}$$

or equivalently precision

$$\begin{aligned} \Big[Var[q_j^2(S_t + \epsilon_{jt})]\Big]^{-1} &= \Big[q_j^4/(1+q_i^2) + q_j^2\Big]^{-1} \\ &= \frac{(1+q_i^2)}{q_j^2(1+q_i^2+q_j^2)} \\ &\triangleq f(q_i, q_j)\,. \end{aligned} \tag{EC.55}$$

Since the precision of prior is $\psi_{ijt}(s)$, it follows that the precision of the posterior is

$$\psi_{ij,t+1}(s) = \psi_{ijt}(s) + \delta(q_i, q_j, a_{ijt}), \tag{EC.56}$$

where $\delta(q_i, q_j, a_{ijt}) \triangleq f(q_i, q_j)a_{ijt}$, which completes the proof for part (i).

To prove part (ii), we use induction on t. First, for $t = 1$ observe that the result holds, because from part (i) we have

$$\begin{aligned} \psi_{ij,2}(s) &= \psi_{ij1}(s) + f(q_i, q_j)a_{ij1} \\ &= \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j)a_{ij1}, \end{aligned} \tag{EC.57}$$

where the second equality follows from Equation (12). Next, assume the required result holds for some $t > 1$:

$$\psi_{ijt}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j)\sum_{l=1}^{t-1} a_{ijl}. \tag{EC.58}$$

We show that it also holds for $t + 1$. To observe this, note that from part (i) we have

$$\psi_{ij,t+1}(s) = \psi_{ijt}(s) + f(q_i, q_j)a_{ijt}. \tag{EC.59}$$

replacing $\psi_{ijt}(s)$ from (EC.58) in (EC.60), we have:

$$\psi_{ij,t+1}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j)\sum_{l=1}^{t-1} a_{ijl} + f(q_i, q_j)a_{ijt} \tag{EC.60}$$

$$= \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j)\sum_{l=1}^{t} a_{ijl}. \tag{EC.61}$$

Thus, the result holds for $t + 1$, which completes the induction step.

Method 2: First observe the following general result (see, e.g., Bishop 2016, page 93[17]). Suppose $\mathbf{X}$ follows a multivariate normal distribution with mean $\boldsymbol{\mu}$ and precision $\boldsymbol{\Lambda}$. That is,

$$\mathbf{X} \sim N(\boldsymbol{\mu}, \boldsymbol{\Lambda}^{-1}).$$

[17] Bishop, C.M. 2016. *Pattern Recognition and Machine Learning*, Springer, New York, NY.

If an observation $\mathbf{Y}$ is made about $\mathbf{X}$, where

$$\mathbf{Y}|\mathbf{X} \sim N(\mathbf{AX}+\mathbf{b},\mathbf{L}^{-1}),$$

then the posterior distribution of $\mathbf{X}$ given $\mathbf{Y}$ is normal and we have

$$\mathbf{X}|\mathbf{Y} \sim N(\mathbf{\Sigma}[\mathbf{A}^T\mathbf{L}(\mathbf{Y}-\mathbf{b})+\mathbf{\Lambda}\boldsymbol{\mu}],\mathbf{\Sigma}),$$

where

$$\mathbf{\Sigma}^{-1} = \mathbf{\Lambda}+\mathbf{A}^T\mathbf{LA}. \tag{EC.62}$$

In particular, upon observing $\mathbf{Y}$, our posterior precision about $\mathbf{X}$ improves by $\mathbf{A}^T\mathbf{LA}$. Now, let $\mathbf{X}=(\alpha_{ijt},\beta_{ijt})^T$ (a column vector), and suppose that at time t it has a bivariate normal distribution with precision $\mathbf{\Lambda}=\mathbf{\Sigma}_{ijt}^{-1}$. Let $\mathbf{A}=(1,s_{t-1})$ and $\mathbf{b}=\mathbb{E}\big[q_j^2(S_t+\epsilon_{jt})\big]$. If sensor i connects to sensor j, then it receives the following signal about $\mathbf{X}=(\alpha_{ijt},\beta_{ijt})^T$:

$$\mathbf{Y}=\alpha_{ijt}+\beta_{ijt}s_{t-1}+q_j^2(S_t+\epsilon_{jt})=\mathbf{AX}+q_j^2(S_t+\epsilon_{jt}).$$

Hence, we have

$$\mathbf{Y}|\mathbf{X} \sim N(\mathbf{AX}+\mathbf{b},\mathbf{L}^{-1}),$$

where

$$\mathbf{L}=\frac{1}{Var[\mathbf{Y}|\mathbf{X}]}=\frac{1}{Var[q_j^2(S_t+\epsilon_{jt})]}=f(q_i,q_j)$$

based on (EC.55). It follows from the general result presented above that the posterior $\mathbf{X}|\mathbf{Y}$ has a bivariate normal distribution. If we denote the precision of this posterior distribution by $\mathbf{\Sigma}_{ij,t+1}^{-1}$, it follows from (EC.62) that

$$\mathbf{\Sigma}_{ij,t+1}^{-1}=\mathbf{\Sigma}_{ijt}^{-1}+\mathbf{A}^T\mathbf{A}f(q_i,q_j)a_{ijt}, \tag{EC.63}$$

where, as before, a_{ijt} is equal to one if i connects to j and is zero otherwise. Equation (EC.63) shows that sensors i's precision about $(\alpha_{ij,t+1},\beta_{ij,t+1})^T$ improves by $\mathbf{A}^T\mathbf{A}f(q_i,q_j)a_{ijt}$ (compared to that of the prior $(\hat{\alpha}_{ijt},\hat{\beta}_{ijt})^T$).

Next, since the posterior $(\hat{\alpha}_{ij,t+1},\hat{\beta}_{ij,t+1})^T$ has a normal distribution with precision given by (EC.63), $\Xi_{ijt+1}(s_{t-1})=\mathbf{A}(\hat{\alpha}_{ij,t+1},\hat{\beta}_{ij,t+1})^T$ is also normally distributed, and its precision defined as $\psi_{ij,t+1}(s_{t-1})$ is given by

$$\psi_{ij,t+1}(s_{t-1})=\big(\mathbf{A}\mathbf{\Sigma}_{ij,t+1}\mathbf{A}^T\big)^{-1}. \tag{EC.64}$$

Note that since the posterior is obtained after receiving a signal but prior to moving to the next period (state change), $\psi_{ij,t+1}(\cdot)$ in (EC.64) is still evaluated at s_{t-1}. Now, using (EC.63) in (EC.64) we have:

$$\begin{aligned}\psi_{ij,t+1}(s_{t-1}) &= \Big(\mathbf{A}\big(\mathbf{\Sigma}_{ijt}^{-1}+\mathbf{A}^T\mathbf{A}f(q_i,q_j)a_{ijt}\big)^{-1}\mathbf{A}^T\Big)^{-1}\\ &= \Big(\mathbf{A}\big(\mathbf{\Sigma}_{ijt}-\frac{f(q_i,q_j)a_{ijt}}{1+g}\mathbf{\Sigma}_{ijt}\mathbf{A}^T\mathbf{A}\mathbf{\Sigma}_{ijt}\big)\mathbf{A}^T\Big)^{-1},\end{aligned} \tag{EC.65}$$

where $g=trace(f(q_i,q_j)a_{ijt}\mathbf{A}^T\mathbf{A}\mathbf{\Sigma}_{ijt})=f(q_i,q_j)a_{ijt}\big(\mathbf{A}\mathbf{\Sigma}_{ijt}\mathbf{A}^T\big)$, and the second equality is due to the following result on inverse of sum of matrices by Miller 1981[18].

[18] K.S. Miller, On the Inverse of the Sum of Matrices, *Mathematics Magazine* 54 (2) (Mar., 1981), 67–72.

Lemma EC.4 **(Miller 1981)**. *If G and $G+H$ are two invertible matrices, and H has rank 1, then let $g = trace(HG^{-1})$. We have that $g \neq -1$ and*

$$(G+H)^{-1} = G^{-1} - \frac{1}{1+g} G^{-1} H G^{-1}.$$

Replacing the value of $g = trace(f(q_i,q_j)a_{ijt}\mathbf{A}^T\mathbf{A}\mathbf{\Sigma}_{ijt}) = f(q_i,q_j)a_{ijt}\big(\mathbf{A}\mathbf{\Sigma}_{ijt}\mathbf{A}^T\big)$ in (EC.65), we have:

$$\begin{aligned}\psi_{ij,t+1}(s_{t-1}) &= \left(\mathbf{A}\mathbf{\Sigma}_{ijt}\mathbf{A}^T\right)^{-1} + f(q_i,q_j)a_{ijt} \\ &= \psi_{ij,t}(s_{t-1}) + f(q_i,q_j)a_{ijt}.\end{aligned} \tag{EC.66}$$

That is, the improvement in the precession $\psi_{ij,t+1}(s_{t-1})$ compared to $\psi_{ijt}(s_{t-1})$ is $f(q_i,q_j)a_{ijt}$, which is independent of s_{t-1}. Finally, from (EC.66) it can be seen that for a general value of s we have

$$\psi_{ij,t+1}(s) = \psi_{ijt}(s) + \delta(q_i,q_j,a_{ijt}), \tag{EC.67}$$

where $\delta(q_i,q_j,a_{ijt}) \triangleq f(q_i,q_j)a_{ijt}$, which is equivalent to (EC.56) in Method 1, and completes the proof for part (i). The proof of part (ii) is the same as that of Method 1. □

Proof of Proposition 3: From Equation (17), we have that a sensor i strictly prefers to target a lower quality sensor (l) than a higher quality one (h) at $t=1$ if, and only if,

$$c_0 + c_1 s_0^2 > 0\,. \tag{EC.68}$$

Hence, it remains to drive sufficient and necessary conditions for the quadratic inequality (EC.68). Doing so, we have that (EC.68) holds, if, and only if, one of the following conditions holds:
(i) $c_0 \leq 0$, $c_1 > 0$, and $|s_0| > \sqrt{-c_0/c_1}$,
(ii) $c_0 > 0$ and $c_1 < 0$, and $|s_0| \leq \sqrt{-c_0/c_1}$, or
(iii) $c_0 > 0$ and $c_1 \geq 0$.
These are conditions (i), (ii), and (iii) of the proposition, and hence, the proof is complete. □

Proof of Lemma 1: Define

$$g_t(s_{t-1}) \triangleq \frac{q_m^2 + 1/\psi_{imt}(s_{t-1})}{q_m^4} - \frac{q_n^2 + 1/\psi_{int}(s_{t-1})}{q_n^4}, \tag{EC.69}$$

and note that by Definition 1 and Equation (10), $m \succeq_{it} n$ if, and only if, $g_t(S_{t-1}) \leq 0$ almost surely. Let

$$f_m \triangleq \sum_{l=1}^{t} a_{iml}, \tag{EC.70}$$

$$f_n \triangleq \sum_{l=1}^{t} a_{inl}, \tag{EC.71}$$

and

$$\Delta f \triangleq f_m - f_n. \tag{EC.72}$$

Next, note that $g_t(S_{t-1}) \leq 0$ almost surely, if, and only if,

$$\Delta f \geq \underline{\Delta} f,$$

where $\underline{\Delta f}$ is some fixed threshold (i.e., a time-independent constant). It remains to show that the inequality $\Delta f \geq \underline{\Delta f}$ is preserved over time (i.e., if it holds at time t then it will hold at any time $t' > t$). To observe this, note that if, at time t, $\Delta f \geq \underline{\Delta f}$, because sensor i sects its target based on

$$\begin{aligned} j^*_{it} &\triangleq \arg\min_{j\in\mathcal{N}\setminus\{i\}} \sigma^2_t(i,j,s_{t-1}) \\ &= \arg\min_{j\in\mathcal{N}\setminus\{i\}} \left\{\frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4}\right\}, \end{aligned} \tag{EC.73}$$

then i will not target n (i.e., $j^*_{it} \neq n$) but may target m (i.e., $j^*_{it} = n$) at time t. In addition, from Proposition 2, we have that

$$\psi_{im,t+1}(s) = \frac{v^2_{im0} w^2_{im0}}{w^2_{im0} + v^2_{im0} s^2} + f(q_i, q_j) \sum_{l=1}^{t} a_{iml}, \tag{EC.74}$$

and

$$\psi_{in,t+1}(s) = \frac{v^2_{in0} w^2_{in0}}{w^2_{in0} + v^2_{in0} s^2} + f(q_i, q_j) \sum_{l=1}^{t} a_{inl}. \tag{EC.75}$$

Hence, Δf at time $t' > t$ will not be less than Δf at time t, if $\Delta f \geq \underline{\Delta f}$ holds at time t. Hence, the inequality $\Delta f \geq \underline{\Delta f}$ is preserved over time. □

<u>Proof of Lemma 2:</u> Note that since

$$\sigma^2_t(i,j,s_{t-1}) = \frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4}, \tag{EC.76}$$

sensor i strictly prefers targeting sensor j to targeting sensor j' at time t if, and only if,

$$\frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4} < \frac{q_{j'}^2 + 1/\psi_{ij't}(s_{t-1})}{q_{j'}^4}. \tag{EC.77}$$

Next, observe that inequality (EC.77) is equivalent to

$$\frac{1}{\psi_{ijt}(s_{t-1})} < \frac{q_j^4}{q_{j'}^4}(q_{j'}^2 + \frac{1}{\psi_{ij't}(s_{t-1})}) - q_j^2 = q_j^2(\frac{q_j^2}{q_{j'}^2} - 1) + \frac{q_j^4}{q_{j'}^4 \psi_{ij't}(s_{t-1})}. \tag{EC.78}$$

Let

$$\bar{\psi}_\epsilon \triangleq [q_j^2(\frac{q_j^2}{q_{j'}^2} - 1)]^{-1}, \tag{EC.79}$$

and note that
(a) $\infty > \bar{\psi}_\epsilon > 0$ because $\frac{q_j}{q_{j'}} > 1 + \epsilon$, and
(b) $\bar{\psi}_\epsilon$ is a fixed number, and is independent of $\psi_{ij't}(s_{t-1})$.
Furthermore,

$$\frac{1}{\bar{\psi}_\epsilon} < q_j^2(\frac{q_j^2}{q_{j'}^2} - 1) + \frac{q_j^4}{q_{j'}^4 \psi_{ij't}(s_{t-1})}, \tag{EC.80}$$

where the inequality (EC.80) follows from (EC.79) and the fact that

$$\frac{q_j^4}{q_{j'}^4 \psi_{ij't}(s_{t-1})} > 0. \tag{EC.81}$$

Therefore, if $\psi_{ijt}(s_{t-1}) > \bar{\psi}_\epsilon$, from (EC.80) we have:

$$\frac{1}{\psi_{ijt}(s_{t-1})} < \frac{1}{\bar{\psi}_\epsilon} < q_j^2(\frac{q_j^2}{q_{j'}^2} - 1) + \frac{q_j^4}{q_{j'}^4 \psi_{ij't}(s_{t-1})}. \tag{EC.82}$$

Hence, if $\psi_{ijt}(s_{t-1}) > \bar{\psi}_\epsilon$, then (EC.78) holds, and therefore, sensor i strictly prefers targeting sensor j to targeting sensor j' at time t. Therefore, $j^*_{it} \neq j'$, and the proof is complete. □

Proof of Proposition 4: We use proof by contradiction. Suppose $\mathcal{T}_i(\mathcal{S}_\infty)$ includes at least two sensors. Since these two sensors differ in their quality by assumption, we can label the higher quality sensor with j and the other sensor with j', and choose $\epsilon > 0$ such that $\frac{q_j}{q_{j'}} > 1+\epsilon$. Note that since $j \in \mathcal{T}_i(\mathcal{S}_\infty)$, sensor i connects to sensor j infinitely often along $\mathcal{S}_\infty$. Therefor, there exists a finite time $\tau(\mathcal{S}_\infty) < \infty$ such that

$$f(q_i, q_j) \sum_{l=1}^{\tau(\mathcal{S}_\infty)} a_{ijl} > \bar{\psi}_\epsilon, \tag{EC.83}$$

where $\bar{\psi}_\epsilon$ is the threshold defined in Lemma 2, and $a_{ijl} = \mathbb{1}_{\{j=j^*_{il}\}}$ along $\mathcal{S}_\infty$. Observe that, since by Proposition 2

$$\psi_{ij,t+1}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j) \sum_{l=1}^{t} a_{ijl}, \tag{EC.84}$$

we have that $\psi_{ijt} > \bar{\psi}_\epsilon$ for all $t > \tau(\mathcal{S}_\infty)$. It then follows from Lemma 2 that $j^*_{it} \neq j'$ for any $t > \tau(\mathcal{S}_\infty)$. This means that i communicates with j' only a finite number of times, which is a contradiction considering the assumption that $j' \in \mathcal{T}_i(\mathcal{S}_\infty)$. □

Proof of Proposition 5: By Proposition 4, $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$ for all $i \in \mathcal{N}$. Let $\kappa_i(\mathcal{S}_\infty)$ denote the unique member of $\mathcal{T}_i(\mathcal{S}_\infty)$. Since sensor i only communicates to $\kappa_i(\mathcal{S}_\infty)$ infinitely often, for all $j \in \mathcal{N} \setminus \{i, \kappa_i(\mathcal{S}_\infty)\}$, there exists a finite time $\tau_{ij}(\mathcal{S}_\infty)$ such that i does not communicate to j after $\tau_{ij}(\mathcal{S}_\infty)$. Let

$$t_i^* = \max_{j \in \mathcal{N} \setminus \{i, \kappa_i(\mathcal{S}_\infty)\}} \tau_{ij}(\mathcal{S}_\infty), \tag{EC.85}$$

and note that $t_i^* < \infty$, since (a) there are a finite number of sensors in $\mathcal{N}$, and (b) all $\tau_{ij}(\mathcal{S}_\infty)$ values are finite. Next, let

$$t^* = \max_{i \in \mathcal{N}} t_i^*, \tag{EC.86}$$

and observe that $t^* < \infty$, since (a) $t_i^* < \infty$ for all $i \in \mathcal{N}$, and (b) $\mathcal{N}$ is a finite set. Moreover, note that at t^* the sets $\mathcal{T}_i(\mathcal{S}_\infty)$ are all determined (for all $i \in \mathcal{N}$). Thus,

$$\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S'}_\infty^{t^*}) \tag{EC.87}$$

for all $i \in \mathcal{N}$, and hence, the proof is complete. □

Proof of Proposition 6: By Proposition 4, $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$ for all $i \in \mathcal{N}$. Let $\kappa_i(\mathcal{S}_\infty)$ denote the unique member of $\mathcal{T}_i(\mathcal{S}_\infty)$. Next, note that from Proposition 5, along any sample path $\mathcal{S}_\infty$, there exists a finite period t^* such that for all $i \in \mathcal{N}$

$$\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S'}_\infty^{t^*}) \quad a.s.$$

Hence, given the sample path up to t^*, we have $\kappa_i(\mathcal{S}_\infty) = \kappa_i(\mathcal{S}'^{t^*}_\infty)$. Let $p_{ij}(\mathcal{S}_\infty|\mathcal{S}_{t^*})$ denote the probability that given the sample path up to t^* (i.e., $\mathcal{S}_{t^*}$), $j \in \mathcal{T}_i(\mathcal{S}_\infty)$. Observe that

$$p_{ij}(\mathcal{S}_\infty|\mathcal{S}_{t^*}) = \sum_{\forall \mathcal{S}_\infty|\mathcal{S}_{t^*}:\ j\in\mathcal{T}_i(\mathcal{S}_\infty)} Pr(\mathcal{S}_\infty|\mathcal{S}_{t^*}) \tag{EC.88}$$

$$= \sum_{\forall \mathcal{S}_\infty|\mathcal{S}_{t^*}:\ j=\kappa_i(\mathcal{S}_\infty)} Pr(\mathcal{S}_\infty|\mathcal{S}_{t^*}) \tag{EC.89}$$

$$= \sum_{\forall \mathcal{S}_\infty|\mathcal{S}_{t^*}:\ j=\kappa_i(\mathcal{S}'^{t^*}_\infty)} Pr(\mathcal{S}_\infty|\mathcal{S}_{t^*}) \tag{EC.90}$$

$$= \mathbb{1}_{\{j=\kappa_i(\mathcal{S}'^{t^*}_\infty)\}}, \tag{EC.91}$$

where the third equality follows from the fact that $\kappa_i(\mathcal{S}_\infty) = \kappa_i(\mathcal{S}'^{t^*}_\infty)$, and the forth equality follows from the fact that $\kappa_i(\mathcal{S}'^{t^*}_\infty)$ is known at time t^*. Hence, we have

$$p_{ij}(\mathcal{S}_\infty|\mathcal{S}_{t^*}) = \begin{cases} 1 : j = \kappa_i(\mathcal{S}'^{t^*}_\infty) \\ 0 : j \neq \kappa_i(\mathcal{S}'^{t^*}_\infty) \end{cases} \tag{EC.92}$$

Therefore, $p_{ij}(\mathcal{S}_\infty|\mathcal{S}_{t^*}) \in \{0,1\}$. □

Lemma EC.5. *A deterministic targeting policy may not be optimal for the unknown-quality sensor targeting problem given by (22)-(24); a randomized policy can be better.*

Proof of Lemma EC.5: Note that under a deterministic policy that prescribes targeting j (almost surely) we can write (24) as:

$$Pr\Big\{V_j \le y_\epsilon\Big\} \ge 1-\epsilon \tag{EC.93}$$

where $V_j \triangleq \Sigma^2(i,j,s_{t-1},Q_j)$ is a random variable with realization $\sigma^2(i,j,s_{t-1})$ (defined in (10)), and Q_j is a random variable with realization q_j.[19] That is, sensor i solves a robust counter part to the original variance reduction problem (see, e.g., (11)) in which instead of connecting to the sensor j that has the minimum $\sigma^2(i,j,s_{t-1})$ it connects to the sensor j that has the minimum

$$F^{-1}_{V_j}(1-\epsilon), \tag{EC.94}$$

where for any random variable Ξ with possible realizations in $\mathcal{Z}$ and c.d.f. F_Ξ

$$F^{-1}_{\Xi}(y) \triangleq \inf\{z\in\mathcal{Z} : F_\Xi(z) \ge y\}. \tag{EC.95}$$

Thus, the optimal robust objective function within the deterministic set of policies denoted by $y^*_d(\epsilon)$ is

$$y^*_d(\epsilon) = \min_{j\in\mathcal{N}\setminus\{i\}} F^{-1}_{V_j}(1-\epsilon), \tag{EC.96}$$

and sensor i connects to sensor

$$j^*_{it,d} = \arg\min_{j\in\mathcal{N}\setminus\{i\}} F^{-1}_{V_j}(1-\epsilon). \tag{EC.97}$$

[19] Note that due to the fact that Q_j is a random variable for sensor i, the familiarity value of i to j is also a random variable. Thus, the familiarity value $\psi_{i,j,t+1}(s_t)$ has a realization which can be calculated based on (15) for each $Q_j = q_j$ that belongs to the ambiguity set considered by sensor i.

In contrast, a feasible (but not necessarily optimal) randomized policy that prescribes connecting to sensor j with probability π_j causes an expected squared error of estimation that is a convex combination of variances: $\bar{V} \triangleq \sum_{j\in\mathcal{N}\setminus\{i\}} \pi_j V_j$. Hence, the robust objective function under this feasible randomized policy denoted by $y_r(\epsilon)$ is

$$y_r(\epsilon) = F_{\bar{V}}^{-1}(1-\epsilon). \tag{EC.98}$$

Since, unlike $F_{V_j}^{-1}$, $F_{\bar{V}}^{-1}$ depends on the joint distribution of all Q_j's defined by $\mathbb{P}_i^Q$ (for all sensor j's where $j \in \mathcal{N}\setminus\{i\}$), it can be the case that $y_r(\epsilon) < y_d^*(\epsilon)$. That is, one can find a feasible randomized targeting strategy for sensor i that is strictly better than any deterministic policy. Because the randomized policy that resulted in $y_r(\epsilon)$ is not necessarily optimal, it follows that the optimal policy π_{it}^* defined in (22) may not be deterministic. □

Proof of Proposition 7: If $|\mathcal{N}| = 2$, the result is trivial. So assume $|\mathcal{N}| > 2$. Consider sensor $i \in \mathcal{N}$, and let $\tilde{\pi} \in \Pi_i$ be a specific randomized communication strategy for sensor i that assigns a probability of one to communicating with sensor $\tilde{j}$, and a probability of zero to communicating with all other sensors $j \in \mathcal{N}\setminus\{i,\tilde{j}\}$. Let $\pi \in \Pi_i$ be any feasible randomized strategy for sensor i. Since for all $j \in \mathcal{N}\setminus\{i\}$

$$V_{\tilde{j}} \le_{s.t.} V_j, \tag{EC.99}$$

we have

$$\bar{V} \triangleq \sum_{j\in\mathcal{N}\setminus\{i\}} \pi_j V_j \ge_{s.t.} \tilde{V} \triangleq \sum_{j\in\mathcal{N}\setminus\{i\}} \tilde{\pi}_j V_j. \tag{EC.100}$$

Next, similar to (EC.98), note that the robust objective function under strategy π denoted by $y_r(\epsilon)$ is

$$y_r(\epsilon) = F_{\bar{V}}^{-1}(1-\epsilon), \tag{EC.101}$$

and under $\tilde{\pi}$ denoted by $\tilde{y}_r(\epsilon)$ is

$$\tilde{y}_r(\epsilon) = F_{\tilde{V}}^{-1}(1-\epsilon). \tag{EC.102}$$

Since $\bar{V} \ge_{s.t} \tilde{V}$, we have that $\tilde{y}_r(\epsilon) \le y_r(\epsilon)$ for all $\epsilon \in [0,1]$. Next, observe that since $\tilde{y}_r(\epsilon) \le y_r(\epsilon)$ holds for any feasible strategy $\pi \in \Pi_i$, we have that strategy $\tilde{\pi}$ is an optimal strategy. Hence, the solution to (22) is given by strategy $\tilde{\pi} \in \Pi_i$. □

Appendix H: Proofs of Results from Appendices

PROOFS FOR RESULTS FROM GENERAL CHANNEL CAPACITY APPENDIX

Proof of Proposition EC.1: Recall that each sensor i is selecting its targets based on the optimization program:

$$\min_{\tilde{s}_{it}\in\mathbb{R},\mathbf{a}_{it}\in\{0,1\}^{n-1}} \quad \mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}\left[\tilde{s}_{it}-S_t\right]^2 \qquad \text{(EC.103)}$$
$$s.t.$$
$$0<c\sum_{j\in\mathcal{N}\setminus\{i\}} a_{ijt}\le b.$$

It follows from the proof of Proposition 1 that

$$\tilde{s}_{it}^*=\mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t].$$

Therefore, in (EC.103), we have

$$\begin{aligned}\min_{\tilde{s}_{it}\in\mathbb{R},\mathbf{a}_{it}\in\{0,1\}^{n-1}} \mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}\left[\tilde{s}_{it}-S_t\right]^2 &= \min_{\mathbf{a}_{it}\in\{0,1\}^{n-1}} \mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}\left[\mathbb{E}_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t]-S_t\right]^2\\ &= \min_{\mathbf{a}_{it}\in\{0,1\}^{n-1}} Var_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t],\end{aligned} \qquad \text{(EC.104)}$$

where the last term denotes sensor i's variance of S_t under the posterior distribution $F_{it}^{\mathbf{a}_{it}}$. That is, sensor i chooses its targets based on minimizing the posterior variance. Furthermore, if sensor i communicates in period t to all sensors in an arbitrary set $\tilde{\mathcal{J}}_{it}\subseteq\mathcal{N}$ satisfying $|\tilde{\mathcal{J}}_{it}|=\lfloor k\rfloor$, then the posterior variance is:

$$Var_{S_t\sim F_{it}^{\mathbf{a}_{it}}}[S_t]=\frac{1}{1+q_i^2+\sum_{j\in\tilde{\mathcal{J}}_{it}}[\sigma_t^2(i,j,s_{t-1})]^{-1}}\,, \qquad \text{(EC.105)}$$

in which $1+q_i^2$ is sensor i's precision prior to communication, and $[\sigma_t^2(i,j,s_{t-1})]^{-1}$ is the precision of the signal it receives from sensor j, where as shown in the main body

$$\sigma_t^2(i,j,s_{t-1}) = \frac{q_j^2+1/\psi_{ijt}(s_{t-1})}{q_j^4}. \qquad \text{(EC.106)}$$

It follows that sensor i must choose set $\tilde{\mathcal{J}}_{it}\subseteq\mathcal{N}$ satisfying $|\tilde{\mathcal{J}}_{it}|=\lfloor k\rfloor$ such that $\sum_{j\in\tilde{\mathcal{J}}_{it}}[\sigma_t^2(i,j,s_{t-1})]^{-1}$ is maximized. Denoting the optimal such set by $\mathcal{J}_{it}^*$, we have that $\mathcal{J}_{it}^*$ is the set of sensors j that have the lowest $\lfloor k\rfloor$ value of $\sigma_t^2(i,j,s_{t-1})$ defined in (EC.106), which means that the proof is complete. □

Proof of Proposition EC.2: The proof directly follows the proof of Proposition 2. □

Proof of Lemma EC.1: The proof follows from the proof of Proposition 2 via pairwise comparison between each member of set $\mathcal{J}$ and sensor j'. Specifically, consider an arbitrary member of set $\mathcal{J}$ and denote it by j. It follows from the proof of Proposition 2 that, at time t, sensor i prefers to communicate to sensor j than sensor j'. Repeating this argument for all $j\in\mathcal{J}$ shows that there are at least $|\mathcal{J}|=\lfloor k\rfloor$ sensors that are better targets than sensor j' from the perspective of sensor i at time t. Since i is allowed to pick only $\lfloor k\rfloor$ sensors as targets, it will not pick sensor j' at time t. Hence, $a_{ij't}^*=0$. □

Proof of Proposition EC.3: We use proof by contradiction. Suppose $\mathcal{T}_i(\mathcal{S}_\infty)$ includes at least $\lfloor k \rfloor + 1$ sensors. Since these sensors differ in their quality by assumption, we can label the sensor with the lowest quality in $\mathcal{T}_i(\mathcal{S}_\infty)$ with j'. Since $\mathcal{T}_i(\mathcal{S}_\infty)$ is a finite set and j' is the sensor with the lowest quality in this set, we can choose $\epsilon > 0$ such that $\frac{q_j}{q_{j'}} > 1 + \epsilon$ for all $j \in \mathcal{T}_i(\mathcal{S}_\infty) \setminus \{j'\}$. Note that sensor i connects to any sensor $j \in \mathcal{T}_i(\mathcal{S}_\infty) \setminus \{j'\}$ infinitely often along $\mathcal{S}_\infty$. Therefor, there exists a finite time $\tau(\mathcal{S}_\infty) < \infty$ such that for all $j \in \mathcal{T}_i(\mathcal{S}_\infty) \setminus \{j'\}$ we have

$$f(q_i, q_j) \sum_{l=1}^{\tau(\mathcal{S}_\infty)} a_{ijl} > \bar{\psi}_\epsilon, \tag{EC.107}$$

where $\bar{\psi}_\epsilon$ is the threshold defined in Lemma EC.1. Observe that, since by Proposition EC.2

$$\psi_{ij,t+1}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j) \sum_{l=1}^{t} a_{ijl}, \tag{EC.108}$$

for all $j \in \mathcal{T}_i(\mathcal{S}_\infty) \setminus \{j'\}$ we have that $\psi_{ijt} > \bar{\psi}_\epsilon$ for all $t > \tau(\mathcal{S}_\infty)$. It then follows from Lemma EC.1 that $a^*_{ij't} = 0$ for any $t > \tau(\mathcal{S}_\infty)$. This means that i communicates with j' only a finite number of times, which is a contradiction considering the assumption that $j' \in \mathcal{T}_i(\mathcal{S}_\infty)$. $\square$

PROOFS FOR RESULTS FROM STATE REALIZATION GENERALIZATION APPENDIX

Proof of Lemma EC.2: The proof is similar to the proof of Lemma 2. In particular, it follows from (EC.14) that sensor i strictly prefers targeting sensor j to targeting sensor j' in period t if, and only if,

$$\frac{q_j^2 + 1/\tilde{\psi}_{ijt}(s_{t-n})}{q_j^4} < \frac{q_{j'}^2 + 1/\tilde{\psi}_{ij't}(s_{t-n})}{q_{j'}^4}, \tag{EC.109}$$

which is equivalent to

$$\frac{1}{\tilde{\psi}_{ijt}(s_{t-n})} < \frac{q_j^4}{q_{j'}^4}(q_{j'}^2 + \frac{1}{\tilde{\psi}_{ij't}(s_{t-n})}) - q_j^2 = q_j^2(\frac{q_j^2}{q_{j'}^2} - 1) + \frac{q_j^4}{q_{j'}^4 \tilde{\psi}_{ij't}(s_{t-n})}. \tag{EC.110}$$

Let

$$\bar{\psi}_\epsilon \triangleq [q_j^2(\frac{q_j^2}{q_{j'}^2} - 1)]^{-1}, \tag{EC.111}$$

and note that
(a) $\infty > \bar{\psi}_\epsilon > 0$ because $\frac{q_j}{q_{j'}} > 1 + \epsilon$, and
(b) $\bar{\psi}_\epsilon$ is a fixed number, and is independent of $\tilde{\psi}_{ij't}(s_{t-1})$.
Furthermore,

$$\frac{1}{\bar{\psi}_\epsilon} < q_j^2(\frac{q_j^2}{q_{j'}^2} - 1) + \frac{q_j^4}{q_{j'}^4 \tilde{\psi}_{ij't}(s_{t-n})}, \tag{EC.112}$$

where the inequality (EC.112) follows from (EC.111) and the fact that

$$\frac{q_j^4}{q_{j'}^4 \tilde{\psi}_{ij't}(s_{t-n})} > 0. \tag{EC.113}$$

Therefore, if $\tilde{\psi}_{ijt}(s_{t-n}) > \bar{\psi}_\epsilon$, from (EC.112) we have:

$$\frac{1}{\tilde{\psi}_{ijt}(s_{t-n})} < \frac{1}{\bar{\psi}_\epsilon} < q_j^2(\frac{q_j^2}{q_{j'}^2} - 1) + \frac{q_j^4}{q_{j'}^4 \tilde{\psi}_{ij't}(s_{t-n})}. \tag{EC.114}$$

Hence, if $\tilde{\psi}_{ijt}(s_{t-n}) > \bar{\psi}_\epsilon$, then (EC.110) holds, and therefore, sensor i strictly prefers targeting sensor j to targeting sensor j' at time t. Therefore, $j^*_{it} \neq j'$, and the proof is complete. □

Proof of Proposition EC.4: Similar to the proof of Proposition 4, we use proof by contradiction. Suppose $\mathcal{T}_i(\mathcal{S}_\infty)$ includes at least two sensors. Since these two sensors differ in their quality by assumption, we can label the higher quality sensor with j and the other sensor with j', and choose $\epsilon > 0$ such that $\frac{q_j}{q_{j'}} > 1 + \epsilon$. Note that since $j \in \mathcal{T}_i(\mathcal{S}_\infty)$, sensor i connects to sensor j infinitely often along $\mathcal{S}_\infty$. Therefore, it is easy to show that there exists a finite time $\tau(\mathcal{S}_\infty) < \infty$ such that $\tilde{\psi}_{ijt}(s_{t-n}) > \bar{\psi}_\epsilon$ for all $t > \tau(\mathcal{S}_\infty)$. It then follows from Lemma EC.2 that $j^*_{it} \neq j'$ for any $t > \tau(\mathcal{S}_\infty)$. This means that i communicates with j' only a finite number of times, which is a contradiction considering the assumption that $j' \in \mathcal{T}_i(\mathcal{S}_\infty)$. □

Proof of Proposition EC.5: Note that by Proposition EC.4, $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$ for all $i \in \mathcal{N}$. Similar to the proof of Proposition 5, let $\kappa_i(\mathcal{S}_\infty)$ denote the unique member of $\mathcal{T}_i(\mathcal{S}_\infty)$. Since sensor i only communicates to $\kappa_i(\mathcal{S}_\infty)$ infinitely often, for all $j \in \mathcal{N} \setminus \{i, \kappa_i(\mathcal{S}_\infty)\}$, there exists a finite time $\tau_{ij}(\mathcal{S}_\infty)$ such that i does not communicate to j after $\tau_{ij}(\mathcal{S}_\infty)$. Let

$$t^*_i = \max_{j \in \mathcal{N} \setminus \{i, \kappa_i(\mathcal{S}_\infty)\}} \tau_{ij}(\mathcal{S}_\infty), \tag{EC.115}$$

and note that $t^*_i < \infty$, since (a) there are a finite number of sensors in $\mathcal{N}$, and (b) all $\tau_{ij}(\mathcal{S}_\infty)$ values are finite. Next, let

$$t^* = \max_{i \in \mathcal{N}} t^*_i, \tag{EC.116}$$

and observe that $t^* < \infty$, since (a) $t^*_i < \infty$ for all $i \in \mathcal{N}$, and (b) $\mathcal{N}$ is a finite set. Moreover, note that at t^* the sets $\mathcal{T}_i(\mathcal{S}_\infty)$ are all determined (for all $i \in \mathcal{N}$). Thus,

$$\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S}'^{t^*}_\infty) \tag{EC.117}$$

for all $i \in \mathcal{N}$, and hence, the proof is complete. □

Proof of Proposition EC.6: The proof follows a similar argument to the proof of Proposition 6. In particular, note that by Proposition EC.4, $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$ for all $i \in \mathcal{N}$. Let $\kappa_i(\mathcal{S}_\infty)$ denote the unique member of $\mathcal{T}_i(\mathcal{S}_\infty)$. Next, note that from Proposition EC.5, along any sample path $\mathcal{S}_\infty$, there exists a finite period t^* such that for all $i \in \mathcal{N}$

$$\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S}'^{t^*}_\infty) \quad a.s.$$

Hence, given the sample path up to t^*, we have $\kappa_i(\mathcal{S}_\infty) = \kappa_i(\mathcal{S}'^{t^*}_\infty)$. Let $p_{ij}(\mathcal{S}_\infty|\mathcal{S}_{t^*})$ denote the probability that given the sample path up to t^* (i.e., $\mathcal{S}_{t^*}$), $j \in \mathcal{T}_i(\mathcal{S}_\infty)$. Observe that

$$p_{ij}(\mathcal{S}_\infty|\mathcal{S}_{t^*}) = \sum_{\forall \mathcal{S}_\infty|\mathcal{S}_{t^*}:\ j \in \mathcal{T}_i(\mathcal{S}_\infty)} Pr(\mathcal{S}_\infty|\mathcal{S}_{t^*}) \tag{EC.118}$$

$$= \sum_{\forall \mathcal{S}_\infty|\mathcal{S}_{t^*}:\ j = \kappa_i(\mathcal{S}_\infty)} Pr(\mathcal{S}_\infty|\mathcal{S}_{t^*}) \tag{EC.119}$$

$$= \sum_{\forall \mathcal{S}_\infty|\mathcal{S}_{t^*}:\ j = \kappa_i(\mathcal{S}'^{t^*}_\infty)} Pr(\mathcal{S}_\infty|\mathcal{S}_{t^*}) \tag{EC.120}$$

$$= \mathbb{1}_{\{j = \kappa_i(\mathcal{S}'^{t^*}_\infty)\}}, \tag{EC.121}$$

where the third equality follows from the fact that $\kappa_i(\mathcal{S}_\infty) = \kappa_i(\mathcal{S}'^{t^*}_\infty)$, and the forth equality follows from the fact that $\kappa_i(\mathcal{S}'^{t^*}_\infty)$ is known at time t^*. Hence, we have

$$p_{ij}(\mathcal{S}_\infty|\mathcal{S}_{t^*}) = \begin{cases} 1 : j = \kappa_i(\mathcal{S}'^{t^*}_\infty) \\ 0 : j \neq \kappa_i(\mathcal{S}'^{t^*}_\infty) \end{cases} \tag{EC.122}$$

Therefore, $p_{ij}(\mathcal{S}_\infty|\mathcal{S}_{t^*}) \in \{0,1\}$. □

PROOFS FOR RESULTS FROM CORRELATED SENSOR ERROR APPENDIX

Proof of Proposition EC.7: It follows from the similar line of proof as Proposition 1 that sensor i will choose it target in period t by selecting the sensor that provides it with the lowest variance of signal. Since when $\nu \neq 0$ the variance of signal i receives from sensor j is given by (EC.30), we have:

$$\begin{aligned} j'^*_{it} &= \operatorname*{argmin}_{j\neq i} \sigma'^2_t(i,j,s_{t-1}) \\ &= \operatorname*{argmin}_{j\neq i} \sigma^2_t(i,j,s_{t-1}) + \frac{\nu^2}{q_i^2} \\ &= \operatorname*{argmin}_{j\neq i} \sigma^2_t(i,j,s_{t-1}) \\ &= j^*_{it}, \end{aligned} \tag{EC.123}$$

where the second line follows from (EC.30), and the third and forth lines follow from Proposition 1. □

Proof of Proposition EC.8: Note that from Proposition 1 we have $j^*_{it} = \operatorname{argmin}_{j\neq i} \sigma^2_t(i,j,s_{t-1})$. Similarly, we have $j'^*_{it} = \operatorname{argmin}_{j\neq i} \sigma'^2_t(i,j,s_{t-1})$, where

$$\sigma'^2_t(i,j,s_{t-1}) = \sigma^2_t(i,j,s_{t-1}) + \frac{\nu^2_{ij}}{q_i^2}. \tag{EC.124}$$

We use proof by contradiction. Suppose $j^*_{it} \neq j'^*_{it}$ and define

$$k_t = \big(\min_{j\neq i, j^*_{it}} \sigma^2_t(i,j,s_{t-1})\big) - \sigma^2_t(i,j^*_{it},s_{t-1}).$$

Note that since $j^*_{it} = \operatorname{argmin}_{j\neq i} \sigma^2_t(i,j,s_{t-1})$, we have $k_t > 0$. Observe that if

$$\max_{j,j'\neq i} \frac{\nu^2_{ij} - \nu^2_{ij'}}{q_i^2} < k_t,$$

then

$$\sigma^2_t(i,j,s_{t-1}) - \sigma^2_t(i,j^*_{it},s_{t-1}) \geq \big(\min_{j\neq i,j^*_{it}} \sigma^2_t(i,j,s_{t-1})\big) - \sigma^2_t(i,j^*_{it},s_{t-1}) = k_t > \max_{j,j'\neq i} \frac{\nu^2_{ij} - \nu^2_{ij'}}{q_i^2} \geq \frac{\nu^2_{ij^*_{it}} - \nu^2_{ij}}{q_i^2} \quad \forall j \neq i, j^*_{it}.$$

Re-arranging the RHS and LHS of the above inequality, we have:

$$\sigma^2_t(i,j,s_{t-1}) + \frac{\nu^2_{ij}}{q_i^2} > \sigma^2_t(i,j^*_{it},s_{t-1}) + \frac{\nu^2_{ij^*_{it}}}{q_i^2} \quad \forall j \neq i, j^*_{it}. \tag{EC.125}$$

The inequality (EC.125) together with (EC.124) implies that $j^*_{it} = \text{argmin}_j\, \sigma'^2_t(i,j,s_{t-1})$. Since by definition $j'^*_{it} = \text{argmin}_j\, \sigma'^2_t(i,j,s_{t-1})$, we have that $j^*_{it} = j'^*_{it}$, which is a contradiction. □

PROOFS FOR RESULTS FROM NON-MYOPIC BEHAVIOR APPENDIX

Proof of Proposition EC.9: First consider the finite-horizon setting. Fix t and information state $(\mathbf{m}, s)$, and let

$$j^* = \arg\min_{j\in\mathcal{N}\setminus\{i\}} \phi_{ij}(m_j, s).$$

Next, for each $j' \in \mathcal{N}\setminus\{i, j^*\}$ define

$$\bar{\eta}^i_{j't}(\mathbf{m},s) \triangleq \frac{\phi_{ij'}(m_{j'},s) - \phi_{ij^*}(m_{j^*},s)}{|\int V^i_{t-1}(\mathbf{m}+\mathbf{e}_{j'},s')\,dF^i_{s'|s}(s') - \int V^i_{t-1}(\mathbf{m}+\mathbf{e}_{j^*},s')\,dF^i_{s'|s}(s')|},$$

and note that $\bar{\eta}^i_{j't}(\mathbf{m},s)$ is positive (and finite) for all $j' \in \mathcal{N}\setminus\{i,j^*\}$. Since $\mathcal{N}$ is a fine set, it follows that

$$\bar{\eta}^i_t(\mathbf{m},s) \triangleq \min_{j'\in\mathcal{N}\setminus\{i,j^*\}} \bar{\eta}^i_{j't}(\mathbf{m},s)$$

is positive (and finite). Furthermore, it follows from the optimality equation

$$V^i_t(\mathbf{m},s) = \min_{j\in\mathcal{N}\setminus\{i\}} \left\{\phi_{ij}(m_j,s) + \eta \int V^i_{t-1}(\mathbf{m}+\mathbf{e}_j,s')\,dF^i_{s'|s}(s')\right\} \tag{EC.126}$$

that for all $\eta \le \bar{\eta}^i_t(\mathbf{m},s)$, it is optimal for sensor i to choose j^* (i.e., the sensor that has the minimum immediate "cost") when there are t periods to go and the information state is $(\mathbf{m},s)$. Finally, since both the number of periods and the number of sensors are finite, we can set

$$\bar{\eta}(\mathbf{m},s) \triangleq \min_{t=1,2,\cdots,T;\, i\in\mathcal{N}} \bar{\eta}^i_t(\mathbf{m},s),$$

where $\bar{\eta}(\mathbf{m},s)$ is positive and is such that for all $\eta \le \bar{\eta}(\mathbf{m},s)$, it is optimal for all sensors to act myopically when the information state is $(\mathbf{m},s)$ (for all $t = 1,2,\cdots,T$), which completes the proof for the finite-horizon setting. The result for the infinite-horizon case follows the same line of proof. In particular, since the optimality equation is

$$V^i_\infty(\mathbf{m},s) = \min_{j\in\mathcal{N}\setminus\{i\}} \left\{\phi_{ij}(m_j,s) + \eta \int V^i_\infty(\mathbf{m}+\mathbf{e}_j,s')\,dF^i_{s'|s}(s')\right\}, \tag{EC.127}$$

it can be seen that fixing the information state $(\mathbf{m},s)$, letting

$$j^* = \arg\min_{j\in\mathcal{N}\setminus\{i\}} \phi_{ij}(m_j,s),$$

and, for each $j' \in \mathcal{N}\setminus\{i,j^*\}$, defining

$$\bar{\eta}^i_{j'\infty}(\mathbf{m},s) \triangleq \frac{\phi_{ij'}(m_{j'},s) - \phi_{ij^*}(m_{j^*},s)}{|\int V^i_\infty(\mathbf{m}+\mathbf{e}_{j'},s')\,dF^i_{s'|s}(s') - \int V^i_\infty(\mathbf{m}+\mathbf{e}_{j^*},s')\,dF^i_{s'|s}(s')|},$$

it is optimal for sensor i to choose j^* (i.e., the sensor that has the minimum immediate "cost") when the information state is $(\mathbf{m},s)$ so long as

$$\eta \le \bar{\eta}^i_\infty(\mathbf{m},s) \triangleq \min_{j'\in\mathcal{N}\setminus\{i,j^*\}} \bar{\eta}^i_{j'\infty}(\mathbf{m},s).$$

Since the number of sensors are finite, we can set $\bar{\eta}(\mathbf{m},s) \triangleq \min_{i\in\mathcal{N}} \bar{\eta}^i_\infty(\mathbf{m},s)$, which completes the proof after nothing that $\bar{\eta}(\mathbf{m},s) > 0$ (because each $\bar{\eta}^i_\infty(\mathbf{m},s)$ is positive). $\square$

Proof of Proposition EC.10: We first show that the value function of each sensor i in both finite-horizon and infinite-horizon settings are non-increasing in $m_{j'}$ for all $j' \in \mathcal{N} \setminus \{i\}$. To this end, let

$$\Delta^i_{j't}(\mathbf{m},s) = V^i_t(\mathbf{m}+\mathbf{e}_{j'},s) - V^i_t(\mathbf{m},s).$$

We use induction to show that for all $t = 1,2,\cdots,T$, $\Delta^i_{j't}(\mathbf{m},s) \le 0$ for all information states $(\mathbf{m},s)$. Since $V^i_0(\mathbf{m},s) = 0$ by definition, from the optimality equation we have

$$V^i_1(\mathbf{m},s) = \min_{j\in\mathcal{N}\setminus\{i\}} \Big\{\phi_{ij}(m_j,s)\Big\}.$$

Thus, it follows from $\phi_{ij'}(m_{j'}+1,s) < \phi_{ij'}(m_{j'},s)$ that $\Delta^i_{j'1}(\mathbf{m},s) \le 0$ for all information states $(\mathbf{m},s)$. Next, suppose $\Delta^i_{j't}(\mathbf{m},s) \le 0$ for some $t = k$ for all information states $(\mathbf{m},s)$. We show that $\Delta^i_{j't}(\mathbf{m},s) \le 0$ for $t = k+1$ for all information states $(\mathbf{m},s)$. From the optimality equation, we have

$$\Delta^i_{j',k+1}(\mathbf{m},s) = \min_{j\in\mathcal{N}\setminus\{i\}} \Big\{\phi_{ij}(m_j + \mathbb{1}_{\{j=j'\}},s) + \eta \int V^i_k(\mathbf{m}+\mathbf{e}_j+\mathbf{e}_{j'},s')\, dF^i_{s'|s}(s')\Big\} \tag{EC.128}$$

$$- \min_{j\in\mathcal{N}\setminus\{i\}} \Big\{\phi_{ij}(m_j,s) + \eta \int V^i_k(\mathbf{m}+\mathbf{e}_j,s')\, dF^i_{s'|s}(s')\Big\} \tag{EC.129}$$

$$\le 0, \tag{EC.130}$$

where the inequality follows from the fact that every term in the minimization in (EC.128) is less than or equal to that in (EC.129) (because of the induction assumption and the fact that $\phi_{ij'}(m_{j'}+1,s) < \phi_{ij'}(m_{j'},s)$). Hence, the finite-horizon value function $V^i_t(\mathbf{m},s)$ is non-increasing in $m_{j'}$ for all $j' \in \mathcal{N} \setminus \{i\}$ and all $t = 1,2,\cdots T$. Taking the limit as $t \to \infty$ then proves that the same property holds for the infinite-horizon value function.

Next, consider the finite-horizon setting, fix t, and suppose it is optimal for sensor i to choose sensor j myopically at some information state $(\mathbf{m}^*,s)$. From the finite-horizon optimality equation, this means that for all $j' \in \mathcal{N} \setminus \{i,j\}$ we have

$$\phi_{ij}(m^*_j,s) \le \phi_{ij'}(m^*_{j'},s)$$

and

$$\phi_{ij}(m^*_j,s) + \eta \int V^i_{t-1}(\mathbf{m}^*+\mathbf{e}_j,s')\, dF^i_{s'|s}(s') \le \phi_{ij'}(m^*_{j'},s) + \eta \int V^i_{t-1}(\mathbf{m}^*+\mathbf{e}_{j'},s')\, dF^i_{s'|s}(s').$$

Since we have $\phi_{ij}(m^*_j+1,s) \le \phi_{ij}(m^*_j,s)$ and $V^i_{t-1}(\mathbf{m}^*+2\mathbf{e}_j,s') \le V^i_{t-1}(\mathbf{m}+\mathbf{e}_j,s')$, it follows from the above inequalities that, for all $j' \in \mathcal{N} \setminus \{i,j\}$, we have

$$\phi_{ij}(m^*_j+1,s) \le \phi_{ij'}(m^*_{j'},s)$$

and

$$\phi_{ij}(m^*_j+1,s) + \eta \int V^i_{t-1}(\mathbf{m}^*+2\mathbf{e}_j,s')\, dF^i_{s'|s}(s') \le \phi_{ij'}(m^*_{j'},s) + \eta \int V^i_{t-1}(\mathbf{m}^*+\mathbf{e}_{j'},s')\, dF^i_{s'|s}(s'),$$

which (based on the optimality equation) means that it is optimal for sensor i to myopically choose sensor j at information state $(\mathbf{m}^*+\mathbf{e}_j,s)$. A similar line of proof shows that the same result holds for the infinite-horizon setting. Finally, if $\bar{m}_j$ denotes the minimum value of m_j such that it is optimal for sensor i to myopically choose sensor j when its number of previous communications with j is $\bar{m}_j$ and with all other sensors $j' \in \mathcal{N} \setminus \{i,j\}$ is $m_{j'} = m^*_{j'}$, then it follows from the previous result that sensor i should myopically choose to communicate to sensor j at any information state $(\mathbf{m}, s)$ with $m_j \geq \bar{m}_j$ and $m_{j'} = m^*_{j'}$ for all $j' \in \mathcal{N} \setminus \{i,j\}$. □

PROOFS FOR RESULTS FROM SELF-UPDATING INFERENCE MODEL APPENDIX

Proof of Lemma EC.3: First consider a specific sensor, sensor i, and for clarity let her precision at time t be denoted by $\mathbf{\Lambda}_i$. Since we have

$$\mathbf{\Sigma}^{-1} = \mathbf{\Lambda}_i + \mathbf{A}^T\mathbf{A}, \tag{EC.131}$$

it follows from Lemma EC.4 (see Online Appendix A) that

$$\mathbf{\Sigma} = \mathbf{\Lambda}_i^{-1} - \frac{1}{1+g}\mathbf{\Lambda}_i^{-1}\mathbf{A}^T\mathbf{A}\mathbf{\Lambda}_i^{-1}, \tag{EC.132}$$

where $g = trace\big(\mathbf{A}^T\mathbf{A}\mathbf{\Lambda}_i^{-1}\big)$. Next, fixing $\epsilon > 0$ and $s = s_{t-1}$, it follows from (EC.132) that we can find $\bar{\mathbf{\Lambda}}_{i,\epsilon,s}$ such that

$$\mathbf{\Sigma}\mathbf{A}^T s_t \leq \frac{\epsilon\mathbf{I} + \mathbf{I}\big[(\alpha,\beta)^T + \mathbf{\Delta}_{it}\big]}{2} \tag{EC.133}$$

and

$$\mathbf{\Sigma}\mathbf{\Lambda}_i\big[(\alpha,\beta)^T + \mathbf{\Delta}_{it}\big] \leq \frac{\epsilon\mathbf{I} + \mathbf{I}\big[(\alpha,\beta)^T + \mathbf{\Delta}_{it}\big]}{2} \tag{EC.134}$$

for all $\mathbf{\Lambda}_i \geq \bar{\mathbf{\Lambda}}_{i,\epsilon,s}$. Adding up (EC.133) and (EC.134), and noting that based on (EC.48) $\mathbf{C}$ for sensor i is given by

$$\mathbf{C} \triangleq \mathbf{\Sigma}\mathbf{A}^T s_t + \big[\mathbf{\Sigma}\mathbf{\Lambda}_i - \mathbf{I}\big]\big[(\alpha,\beta)^T + \mathbf{\Delta}_{it}\big] \tag{EC.135}$$

we have that (for sensor i) $\mathbf{C} \leq \epsilon\mathbf{I}$ when $\mathbf{\Lambda}_i \geq \bar{\mathbf{\Lambda}}_{i,\epsilon,s}$. A similar line of proof shows that (for sensor i) $\mathbf{C} \geq -\epsilon\mathbf{I}$ when $\mathbf{\Lambda}_i \geq \bar{\mathbf{\Lambda}}_{i,\epsilon,s}$. Finally, since the number of sensors is finite we can find $\bar{\mathbf{\Lambda}}_{\epsilon,s}$ such that $\bar{\mathbf{\Lambda}}_{\epsilon,s} \geq \bar{\mathbf{\Lambda}}_{i,\epsilon,s}$ for all $i \in \mathcal{N}$. Using the threshold $\bar{\mathbf{\Lambda}}_{\epsilon,s}$ then completes the proof. □

Proof of Proposition EC.11: Recall that

$$j^*_{it} \triangleq \arg\min_{j \in \mathcal{N}\setminus\{i\}} \sigma^2_t(i,j,s_{t-1}) \tag{EC.136}$$

where

$$\sigma^2_t(i,j,s_{t-1}) = Var\big[\epsilon_{jt} + \frac{\hat{\alpha}_{ijt} + \hat{\beta}_{ijt}s_{t-1}}{q^2_j}\big].$$

Now let j^{**}_{it} be the counterpart to j^*_{it} denoting sensor i's optimal target at time t with the difference that j^{**}_{it} indicates the optimal target when sensors update their inference model parameters. Since the number of sensors are finite, observe from (EC.136) that we can pick an $\epsilon^* > 0$ such that perturbing all $\hat{\alpha}_{ijt}$ and $\hat{\beta}_{ijt}$ parameters (i.e., for all $j \in \mathcal{N} \setminus \{i\}$) by at most $\pm\epsilon^*$ does not change j^*_{it}. Furthermore, if we set $s = s_{t-1}$ and denote the changes in sensor j's parameters due to updating by $\mathbf{C}_j$, it follows from Lemma EC.3 that there exists a threshold $\bar{\mathbf{\Lambda}}_{\epsilon^*,s}$ such that for all $j \in \mathcal{N} \setminus \{i\}$ we have $-\epsilon^*\mathbf{I} \leq \mathbf{C}_j \leq \epsilon^*\mathbf{I}$ if the precision of all sensors is greater than or equal to $\bar{\mathbf{\Lambda}}_{\epsilon^*,s}$. Since when $-\epsilon^*\mathbf{I} \leq \mathbf{C}_j \leq \epsilon^*\mathbf{I}$, each sensor $j \in \mathcal{N} \setminus \{i\}$ will change its parameters by at most ϵ^* and sensor i is aware of this, sensor i will not change its $\hat{\alpha}_{ijt}$ and $\hat{\beta}_{ijt}$ parameters by more than ϵ^* when updating is allowed. Thus, for sensor i,

$$j^*_{it} = j^{**}_{it}$$

so long as the precision of all sensors is greater than or equal to $\bar{\mathbf{\Lambda}}_{\epsilon^*,s}$. Repeating this argument for all other sensors and letting $\bar{\mathbf{\Lambda}}_s$ be the maximum of the thresholds among all sensors completes the proof. □

Appendix I: Numerical Studies Detailed Description and Discussion

Known Qualities

We present the results of two numerical studies that examine, respectively, the roles that the underlying state dynamics and the initial state distribution play in long run communication network formation. In both studies, for ease of illustration, we consider a collection of three sensors, i.e., $\mathcal{N} = \{1,2,3\}$. The following sensor characteristics are adopted for each study. Sensor qualities are asymmetric and decreasing in sensor labels, with $\underline{q} \triangleq (q_j : j \in \mathcal{N}) = (9.8, 9.6., 9.4)$. The initial familiarities v_{ij0} and w_{ij0} are specified by the following matrix:

$$[v_{ij0}]_{i,j\in\mathcal{N}} = [w_{ij0}]_{i,j\in\mathcal{N}} = \begin{pmatrix} \infty & 1.6 & 2.4 \\ 0.8 & \infty & 2.4 \\ 0.8 & 1.6 & \infty \end{pmatrix}. \tag{EC.137}$$

Thus, any two different sensors i and k have the same initial familiarities about the third sensor j. This matrix also implies that initial familiarities are increasing in sensor labels. Combined with the above quality-labeling scheme, it then follows that sensor 1 is the most attractive from a quality perspective but that sensor 3 is the most attractive from an initial familiarity perspective. Sensor 2 lies in between sensors 1 and 3 in that it represents mid values of both quality and familiarity. We emphasize that by design these parameters ensure that no sensor is initially dominant from the perspective of any other sensor. Also, these parameters ensure that target selection in the initial period is given by Proposition 3 (i); that is, the intermediate region of Figure 1 applies and so high (absolute) state values favor familiarity over quality.

For each problem instance in each study, the time-0 random directed graph $\vec{G}(\mathcal{N}, \mathcal{E}, \mathcal{P})$ was generated by simulating $1,000$ sample paths, each representing a distinct realization of the underlying AR(1) process over time. That is, at time $t = 0$, the probability that j will be the eventual long run target of i (see p_{ij} defined in (21)) is estimated using the outcomes of the $1,000$ sample paths.[20]

Study 1 (The Effect of State Dynamics): To capture the effect of state dynamics, we consider nine different cases for the environment's AR(1) parameters (α, β). The nine cases represent pairwise combinations of low, medium, and high values of both α and β: $\alpha \in \{0,2,4\}$ and $\beta \in \{0.3, 0.5, 0.7\}$. The initial state is randomly drawn from a normal distribution with a mean of 10 and a standard deviation of 2 in all cases.

Figure 2 (shown in the main paper) presents the long run communication networks for all nine cases of (α, β). Recall that by design no sensor is dominant from the perspective of any other sensor. Despite this fact, observe in the top left panel (low α and low β) that the long run contact of each sensor is always the high-quality sensor, i.e., 1 eventually always targets 2, and 2 and 3 eventually always target 1. Also, observe in bottom right panel (high α and high β) that the long run contact of each sensor is always the high-familiarity sensor, i.e., 1 and 2 target 3, and 3 targets 2. In contrast, at intermediate values of α and β, e.g., the center panel, the long run contact of each sensor typically depends on the sample path: for some sample paths, the higher quality sensor wins and for others the higher-familiarity sensor wins. However, we observe that the winner is more likely to be the more familiar sensor as either α and β increases. The reason for this (α, β) effect is twofold. One, high state values favor more familiar targets because Proposition 3 (i) applies (by study design). Two, higher state values are more likely to occur on any given sample path as

[20] The number of simulated sample paths were chosen so that the point estimations for p_{ij} values have a low enough standard error, and hence, provide reliable estimation.

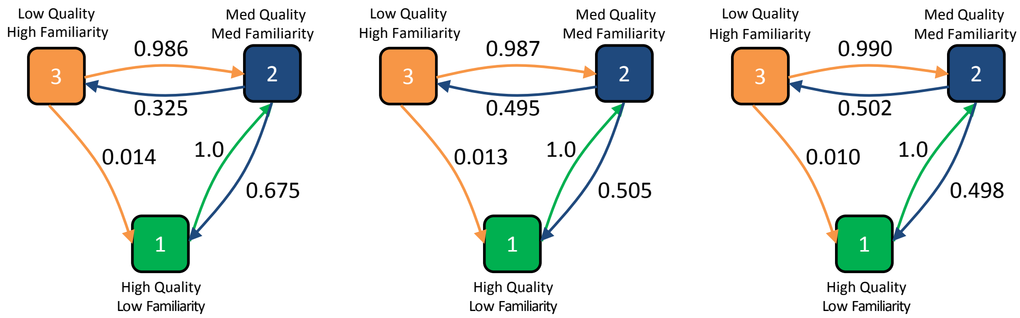

(a) $(\mu,\sigma)=(2,2)$ (b) $(\mu,\sigma)=(10,2)$ (c) $(\mu,\sigma)=(18,2)$

Figure EC.1 The role of initial state in long run communication network. S_0 has a $N(\mu,\sigma)$ distribution.

either α or β increases. Therefore, higher-familiarity sensors are increasingly favored as α and/or β increases. Interestingly, we also note that α typically has a stronger effect than β.[21]

Study 2 (The Effect of Initial State): We now fix the environment's AR(1) parameters as $(\alpha,\beta)=(2,0.5)$ and consider three different $N(\mu,\sigma)$ distributions for the initial state: (a) $(\mu,\sigma)=(2,2)$; (b) $(\mu,\sigma)=(10,2)$, and (c) $(\mu,\sigma)=(18,2)$. We know from our earlier theoretical results that the initial state can impact the long run contact probability through its impact on initial selection and the resulting familiarity gain. That is, due to the stickiness factor, initial state can create a momentum that might last for ever.

The long run communication (targeting) network for each case is presented in Figure EC.1. Observe that as we move from left to right, i.e., from case (a) to (b) to (c), sensors 2 and 3 settle on their higher-familiarity sensor (3 for 2; 2 for 3) with a somewhat higher probability. This reflects the fact that higher initial states favor initial selection of the more familiar sensor and higher initial states are more likely as we move from (a) to (b) to (c). For sensor 1, although there is a higher probability of the more familiar sensor being initially targeted as we move from case (a) to (b) to (c), this has no effect on the long run contact probability because quality eventually wins over initial familiarity with all sample paths.

In summary, the initial state can impact the long run contact probability through its impact on initial selection and the resulting familiarity gain. However, in comparing Study 1 and Study 2, we observe that the effect of the initial state is less strong than the effect of the underlying state dynamics, i.e., the (α,β) parameters. This is because the (α,β) parameters influence the state realization in every period, whereas the initial state's impact on future states can diminish over time.

Unknown Qualities

Next, we numerically explore how the tradeoffs we observed in our earlier results and experiments between (a) quality, (b) familiarity, and (c) state can be affected by the underlying ambiguity around qualities and/or by the level of optimism of sensors.[22] In both of the following studies (Studies 3 and 4 below), we assume that a collection of three sensors operates in an environment where the underlying AR(1) model has parameters $(\alpha,\beta)=(2,0.7)$. We assume that sensor $i\neq j$ believes that sensor j's quality lies in the range $(0,2q_j)$ with all values in this range equally likely and where $\underline{q}\triangleq(q_j:j\in\mathcal{N})=(9.8,9.6.,9.4)$. Thus, the average perceived quality and ambiguity range

[21] We note that at higher values of β, e.g., $\beta=0.9$ (not shown), the higher familiarity sensor has a small probability of being the long run contact even when $\alpha=0$.

[22] We restrict our attention to the set of deterministic policies to gain clear insights, and avoid extra levels of complexities that are caused by randomized policies.

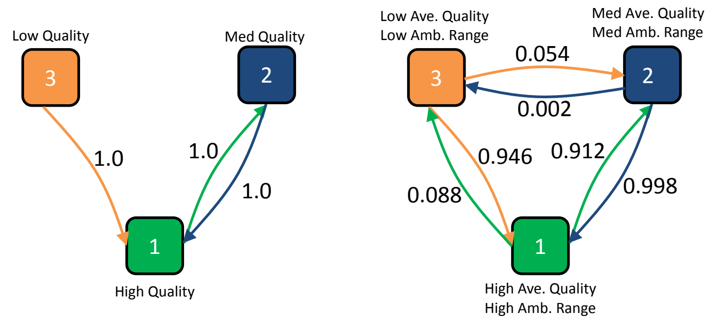

(a) Known Qualities (b) Unknown Qualities

Figure EC.2 With unknown qualities, the communication network is stochastic even with identical initial familiarities.

both decrease in the sensor label, such that Sensor 1 (3) has both the highest (lowest) maximum possible quality and the highest (lowest) average perceived quality. That is, while a higher-label sensor is perceived to have a higher quality on average, a lower-label sensor's quality is known with less ambiguity. In what follows, the known-quality counterpart problem sets the qualities at the exact q_j values given by $\underline{q} \triangleq (q_j : j \in \mathcal{N}) = (9.8, 9.6., 9.4)$, which is (a) similar to values used in our earlier studies, and (b) corresponds to the average perceived qualities in the unknown-quality case.
Study 3 (Common Initial Familiarities) In this study, we assume initial familiarity values are common among sensors (set at $v_{ij0} = w_{ij0} = 2$), and compare the cases of known and unknown qualities. For the case of known qualities, we analytically established earlier that target selection is deterministic (i.e., sample path independent) and time-invariant in the case of common initial familiarities (see earlier Special Case 1): the highest-quality sensor targets the second-highest quality sensor, all other sensors target the highest quality sensor. This is illustrated in Figure EC.2(a) [main paper] in which the link value is 1 on $1 \to 2$; $2 \to 1$; and $3 \to 1$. The long run contact may not be deterministic, however, when qualities are unknown; see, e.g., Figure EC.2(b) [main paper] which presents the long run communication network for an optimism parameter of $\epsilon = 0.2$.

Observe from the link values that in the long run each sensor typically (but not always) targets the lowest-label sensor available to it, indicating that each sensor's long run target is more likely to be the one for which the ambiguity is lowest even though such a sensor's average perceived quality is the lowest. However, this preference towards lower ambiguity (which comes at the cost of targeting a lower average quality sensor) is violated on some sample paths. A value of $\epsilon = 0.2$ implies that sensors are quite pessimistic (i.e., ambiguity-averse) in target selection, and hence they put a lot of weight on avoiding targeting a potentially low-quality sensor. Given our construction of the ambiguity sets, lower-label sensors have lower likelihoods of low quality values, and that is why sensors tend to target low-label sensors. However, because there is some chance that the lower label sensor may be the one with the lower quality (e.g., higher labeled sensors have a higher maximum possible quality), we observe that this long run low-label selection does not occur along all sample paths.
Study 4 (The Effect of the Optimism Level): We now explore the impact of the optimism parameter ϵ on long run target selection. We consider a setting where initial familiarities vary across sensors. In particular, we use the same initial familiarity matrix as used in our earlier known-quality studies; see (EC.137). Therefore, the study replicates the known-quality instance in Figure 2(f) except that there is now quality ambiguity; sensors can't assume that qualities are exactly at their means as was the case with known qualities. Figure 3(a) replicates the known-quality communication network, i.e., Figure 2(f), for ease of comparison. Figures 3(b) and 3(c) present the unknown-quality communication networks for optimism parameters $\epsilon = 0.2$ and $\epsilon = 0.9$,

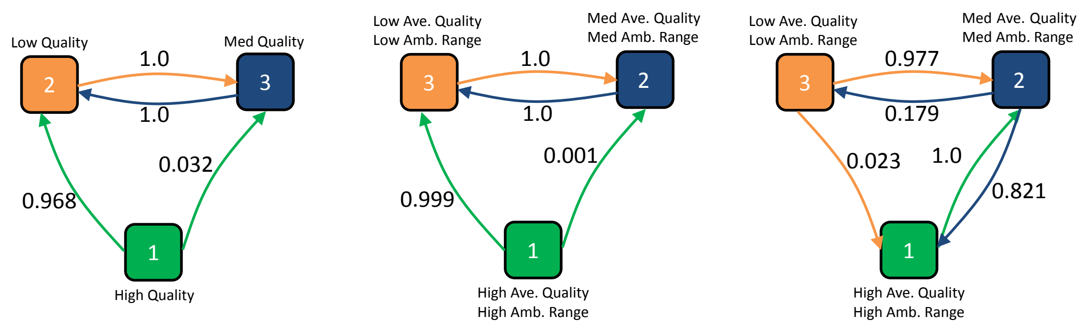

(a) Known Qualities (b) Unknown Qualities, $\epsilon = 0.2$ (c) Unknown Qualities, $\epsilon = 0.9$

Figure EC.3 Communication network differs when qualities are unknown and depends on the optimism parameter ϵ.

.

respectively. One immediate observation is that the link probabilities differ when qualities are unknown. This is because target selection now must consider ambiguities in addition to qualities and familiarities. Moreover, we observe that the optimism parameter strongly influences long run target selection. Sensors are more pessimistic (optimistic) about potential sensor qualities when ϵ is low (high) and this in turn influences the emphasis placed on familiarity versus possible qualities, which in turn influences the role of state in target selection. The optimism parameter used by a sensor will depend on the firm that deployed it. Thus, our results imply that organizational attitudes towards ambiguity will impact target selection and the resulting communication network that evolves over time.

Combining findings from this study and those presented earlier, we see that the inherent targeting trade-off between quality and familiarity is influenced by both the environment (through the state dynamics) and the firms deploying the sensors (through the ambiguity optimism parameter).