MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 24, No. 1, January–February 2022, pp. 333–351
ISSN 1523-4614 (print), ISSN 1526-5498 (online)
http://pubsonline.informs.org/journal/msom

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

Soroush Saghafian,[a] Brian Tomlin,[b] Stephan Biller[c]

[a] Harvard Kennedy School, Harvard University, Cambridge, Massachusetts 02138; [b] Tuck School of Business, Dartmouth College, Hanover, New Hampshire 03755; [c] Advanced Manufacturing International, Clearwater, Florida 33762

Contact: soroush_saghafian@hks.harvard.edu, https://orcid.org/0000-0002-9781-6561 (SS); brian.tomlin@tuck.dartmouth.edu, https://orcid.org/0000-0002-2533-1039 (BT); sbiller@umich.edu (SB)

Received: November 24, 2019
Revised: July 12, 2020
Accepted: October 6, 2020
Published Online in Articles in Advance: March 17, 2021

https://doi.org/10.1287/msom.2020.0958

Abstract. ***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, 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 (1) are constrained in the number of sensors they can target and (2) only have partial knowledge of how other sensors operate—that is, they do not know others' underlying inference algorithms/models. We study the targeting problem, examine the evolution of interfirm 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 among 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 is used. ***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. 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 among which it cycles 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.

Supplemental Material: The online appendix is available at https://doi.org/10.1287/msom.2020.0958.

Keywords: estimation • information sharing • sensor collaboration • robust optimization

1. Introduction

The internet of things (IoT) and its sensor-based monitoring offer 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—for example, 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). This McKinsey (2015) report claims that there is a vast untapped potential to use sensor information to make better operations decisions related to predictive maintenance and worker safety, and the report states that "the real value of IoT applications [comes] from analyzing data from multiple sensors" (McKinsey 2015, 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 on 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—for example, 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 among vibration sensors on different pieces of equipment or among 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 earlier, many thousands of critical equipment assets. Although 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 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). Although estimation performance would be highest if firms shared all relevant information, sensor technology (e.g., inference model, measurement technique, and raw readings) is often proprietary and closely guarded, 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: (1) generating better operating-condition information through sensor communication, and (2) 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 the first aspect—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 the second aspect 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 therefore can 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, pp. 4–5) is envisioned as a remedy for sensor quality concerns in a wide range of IoT settings, including energy, environmental monitoring, healthcare, infrastructure, and various other industrial applications. Furthermore, 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" (International Electrotechnical Commission (IEC) 2014, p. 33). In healthcare, for example, many firms are developing minimally invasive wearable biosensors 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 2020)—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. 2020), industrial firms will increasingly seek to understand the formation and evolution of autonomous interfirm 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—that is, multifirm—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 Online Appendix A, we show how our results extend naturally 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 (1) the qualities of both the targeted and targeting sensors, and (2) the targeting sensor's 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—that is, 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 multifirm 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 helps 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—for example, 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 this paper is organized as follows. The most relevant literature is discussed in Section 2. The base model is described in Section 3. Analysis and results are presented in Sections 4 and 5. The extension to unknown sensor quality is developed and analyzed in Section 6. A number of other extensions are summarized in Section 7 and fully developed in the online appendices. Conclusions are discussed in Section 8. Appendices A–C are available in the online appendix. Appendices D–I are available in an unabridged version (Saghafian et al. 2018) available on the Social Science Research Network (SSRN).

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. Multisensor 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 because of 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 explore 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 2020). Through that lens, one can view our work as exploring a related but different question: when each individual in a 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 this lens in mind, the paper most related to our work appears to be that of 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 trade off between quality and unknown beliefs), our work differs significantly from that of 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 from 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 important, 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. Importantly, ambiguity differs from risk, and we refer to Saghafian (2018) and the references therein for related discussions.

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 (2021), and references therein. For some general theoretical results on the percentile optimization approach that we use, 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—for example, 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 as 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, p. 6) that can then be fed back to the sensor.[3] Alternatively, 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, also assume that it incurs a delay of one period and is perfectly accurate. In Online 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 this 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, target-selection problem whereby each sensor chooses from which other sensors to solicit state estimates.

3.1. 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 independent and identically distributed (i.i.d.) normal white-noise random variables with mean zero and variance normalized to one. 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 level (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—for example, Vitus et al. (2012), Shi and Zhang (2012), and Ballotta et al. (2019).[4]

3.2. 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—for example, 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 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 zero 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 firm'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 preceding 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 preceding 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 one. On realizing the current signal $\Gamma_{it} = \gamma_{it}$, sensor i updates its prior belief about the current state according to the Bayes' rule. Because both the signal received about the state and the prior on state have a normal distribution (see (2) and (3)), it follows from the Bayes' rule that sensor i's posterior belief about the state is also normally distributed but with a mean and variance given by

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

and

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

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

3.3. 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)), 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 (predeployment) 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 Section 6.

3.4. Target Selection

In each period t, after updating its state estimate based on its private signal as in (4) and (5), each sensor i chooses a set of sensors from which to request state estimates—that is, 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 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:

$$\begin{aligned} \min_{\tilde{s}_{it} \in \mathbb{R}, \mathbf{a}_{it} \in \{0,1\}^{n-1}} \quad & \mathbb{E}_{S_t \sim F_{it}^{\mathbf{a}_{it}}}[\tilde{s}_{it} - S_t]^2 \\ \text{subject to (s.t.)} \quad & 0 < c \textstyle\sum_{j \in \mathcal{N} \setminus \{i\}} a_{ijt} \le b, \end{aligned} \tag{6}$$

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 on the vector $\mathbf{a}_{it} \in \{0,1\}^{n-1}$ (to whom 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 sensor is trying to communicate with 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—that is, $\lfloor k \rfloor = 1$—because this is where 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 this paper and fully show in Online Appendix A.

3.5. 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 nonmyopic 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 Section 7, but we relegate their detailed analysis and discussion to the online 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.

4.1. Targeting Equivalence

We begin by establishing that the target-selection problem in (6) 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 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 $\mathbb{E}[S_{jt}|\Gamma_{jt}]$—that is, 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

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

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 later.) 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:

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

We denote the variance in this signal's noise as

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

There are two independent sources of noise in this signal: (1) the inherent white noise ϵ_{jt} in sensor j's measurement Γ_{jt} (which has a variance of $1/q_j^2$), and (2) 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 $\psi_{ijt}(s_{t-1}) \triangleq 1/\mathrm{Var}[\Xi_{ijt}(s_{t-1})]$, 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

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

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.*

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 Online Appendix A for more details and an 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*—that is, 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 establish later, 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 $\psi_{ijt}(s)$.

4.2. 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 $t+1$—that is, how $\psi_{ijt}(s)$ updates to $\psi_{ij,t+1}(s)$. To that end, we first note that it follows from the definition $\psi_{ijt}(s_{t-1}) \triangleq 1/\mathrm{Var}[\Xi_{ijt}(s_{t-1})]$ 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 $\Xi_{ijt}(s_{t-1}) = \hat{\alpha}_{ijt} + \hat{\beta}_{ijt}s_{t-1}$:

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

There are two independent sources of noise in this signal: (1) 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 (2) the inherent white noise ϵ_{jt} in sensor j's measurement, which has a variance of $1/q_j^2$. Thus, based on (13), the variance in the signal's noise is given by $\mathrm{Var}[q_j^2(S_t + \epsilon_{jt})] = q_j^4/(1+q_i^2) + q_j^2$. 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 $s \in \mathbb{R}$:*

a. The familiarity function satisfies $\psi_{ij,t+1}(s) = \psi_{ijt}(s) + \delta(q_i, q_j, a_{ijt})$, *where* $\delta(q_i, q_j, a_{ijt}) \triangleq f(q_i, q_j)a_{ijt}$ *and*

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

b. *For all $t = 1, 2, 3, \cdots$, we have*

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

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—that is, $a_{ijt} = 1$. 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 important, 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 with 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 Whom?

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, besides 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 implicitly 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]

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 this 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, $0 < r \le 1$. Using (11) and (12), it follows that sensor i strictly prefers targeting the lower-quality sensor (l) if and only if

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

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

$$c_0 + c_1 s_0^2 > 0, \qquad (17)$$

where

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

and

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

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.

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*:

a. $c_0 \le 0$, $c_1 > 0$, *and* $|s_0| > \sqrt{-c_0/c_1}$,
b. $c_0 > 0$ *and* $c_1 < 0$, *and* $|s_0| \le \sqrt{-c_0/c_1}$, *or*
c. $c_0 > 0$ *and* $c_1 \ge 0$.

This proposition highlights the interconnected roles that (1) sensor qualities, (2) familiarities, and (3) state play in target selection. Intuitively, if sensor i has more familiarity with the high-quality sensor's inference model parameters (i.e., $v_{ih0} \ge v_{il0}$ and $w_{ih0} \ge w_{il0}$), then the high-quality sensor is the inherently more attractive target regardless of state. This is reflected in Proposition 3 by the fact that $c_0 < 0$ and $c_1 < 0$ in this case, and therefore, sensor h is preferred. By contrast, 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 (a) and (b) of Proposition 3) unless the familiarity advantage of the lower-quality sensor compared with the higher-quality one is so large that it makes the lower-quality sensor the preferred target regardless of the state (part (3) of Proposition 3).

As the quality ratio r increases from zero to one (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—that is, sensor h is of much higher quality than l—then h dominates l—that is, h is targeted in all states. This h-dominating region always exists, but it does not cover the entire range $0 < r \le 1$ unless $v_{ih0} \ge v_{il0}$ and $w_{ih0} \ge w_{il0}$—that is, when the high-quality sensor is more familiar for both parameters. In contrast, when the quality ratio is high—that is, sensor qualities are relatively similar—then l dominates h, meaning that l is targeted in all states. This l-dominating region exists if and only if $v_{ih0} < v_{il0}$ and $w_{ih0} < w_{il0}$—that is, when 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(a) 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 (b) applies. When this intermediate region exits, then case (a)—that is, high state favors high familiarity sensor—applies over this entire intermediate region if $r > \sqrt{w_{ih0}/w_{il0}}$, but case (b)—that is, 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,

Figure 1. (Color online) Initial Selection Between a Higher-Quality Sensor (h) and a Lower-Quality Sensor (l)

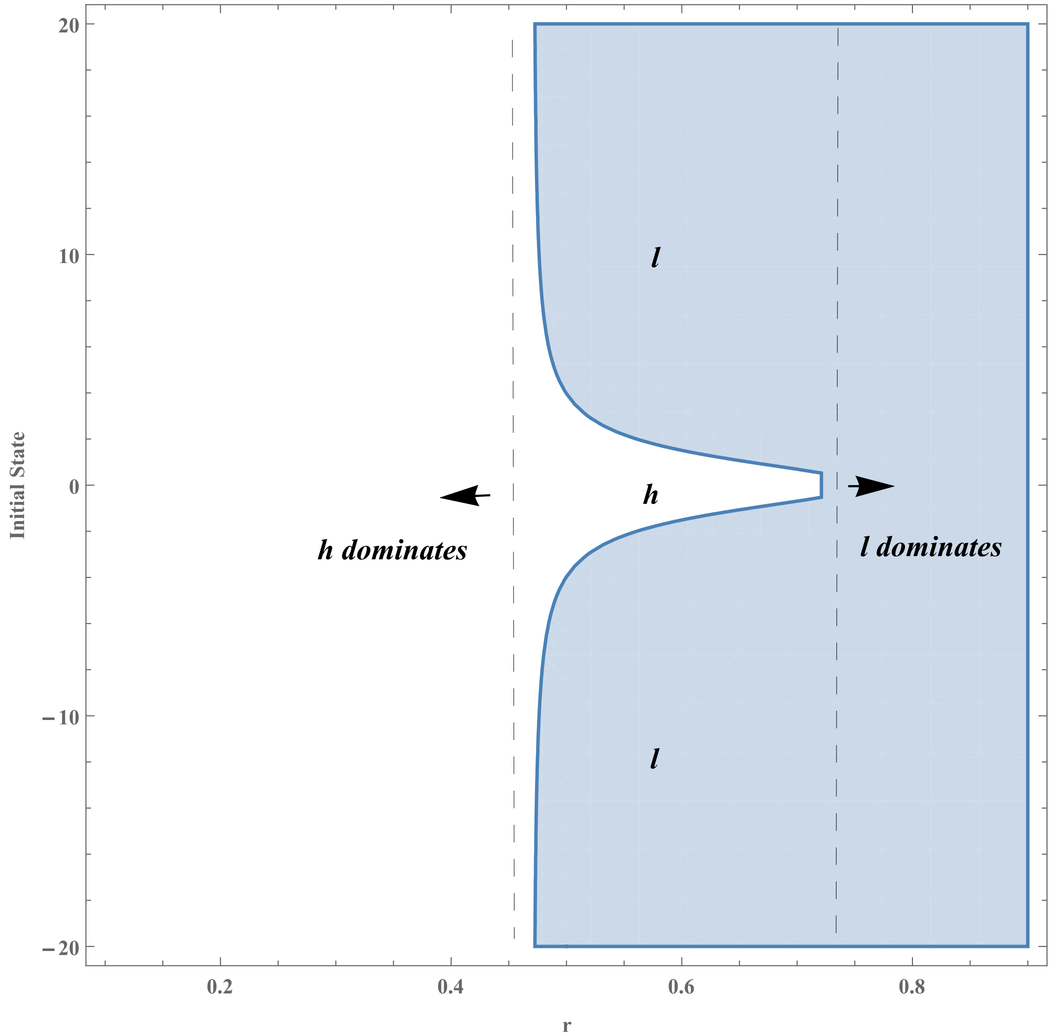

Note. The darker (lighter) area represents the region in which l (h) is the preferred sensor (Proposition 3(a) applies in the intermediate region between the dashed lines).

we turn our attention to exploring how target selection evolves over time.

5.2. Target Evolution and Long-Run Target Selection

Analogous 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 because of 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 now it also depends on (1) i's communication history with j as reflected by the number of times i targeted j in the past, and (2) 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, so the attractiveness of j as a future target for i increases every 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—that is, 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), 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, 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, 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, (1) higher-quality sensors can build familiarity with other sensors more rapidly than can lower-quality sensors, and (2) 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 $\Pr(\sigma_t^2(i, m, S_{t-1}) \leq \sigma_t^2(i, n, S_{t-1}) | \mathcal{H}_t) = 1$, 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 $j \in \mathcal{N} \setminus \{i\}$—that is, $v_{ij0} = v_{i0}$ 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 of any other sensor's parameters. It follows from part (c) 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 heterogeneous 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—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 $\mathcal{S}_\infty'^t$ denote a sample path that is equivalent to $\mathcal{S}_\infty$ up to time t but one that may deviate from $\mathcal{S}_\infty$ afterward: $\mathcal{S}_\infty'^t \triangleq \mathcal{S}_t \cup \{s'_{t+1}, s'_{t+2}, \cdots\}$. 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 $\mathcal{S}_\infty$, (1) $\mathcal{T}_i(\mathcal{S}_\infty)$ is a singleton—that is, $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$—and (2) the unique long-run contact in $\mathcal{T}_i(\mathcal{S}_\infty)$ can be identified in the almost sure sense in finite time—that is, $\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S}_\infty'^{t^*})$ *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*—that is, 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—that is, 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 $j^*_{it} \neq j'$.*

This lemma states that a sensor j' will not be targeted by a sensor i if (1) there is another sensor j of a higher quality than j', and (2) 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, Lemma 2 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 later), 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 (1) sample-path dependent, and (2) 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. Because of 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. Because dominance persists (see Lemma 1), the higher-quality sensor may not be the long-run contact of sensor i.

Using the preceding 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—that is, 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 $i \in \mathcal{N}$,*

$$\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i\left(\mathcal{S}_\infty^{*t^*}\right) \quad \textit{almost surely.}$$

This 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 earlier with the additional property that $p_{ij} \in \{0, 1\}$, for all $i, j \in \mathcal{N}$.*

In Online Appendix A, we show that all the preceding 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.

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 earlier allow a firm to predict which other firms will have a high likelihood of being its long-run information suppliers, both for (1) an individual asset on a particular worksite, and (2) 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—for example, 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 its suppliers. For example, for reasons of corporate relationships or geopolitical considerations, a firm might be more comfortable being dependent on some firms rather than others. From our results, 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 the same manner as their sensors), if firm i would prefer to be dependent on firm j over firm 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 other's sensors are more likely to end up as each other's targets, 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 contacts of other firms' sensors and, therefore, to be more connected to other firms—that is, be a central node. The implication of this observation is that firms with high-quality sensors will serve as information suppliers to many other 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, which 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 midvalues of both quality and initial familiarity. (Full details and discussion of this numeric study and others can be found in Appendix I of the unbridged version (Saghafian et al. 2018).)

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—that is, sensor 1 eventually always targets sensor 2, and sensors 2 and 3 eventually always target sensor 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— that is, sensors 1 and 2 target sensor 3, and sensor 3 targets sensor 2. In contrast, at intermediate values of α and β—for example, in 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 earlier) influence its centrality in the pattern irrespective of the underlying dynamics through its sensor quality and its 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. This 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.

Figure 2. (Color online) The Influence of Underlying State Dynamics on Communication Pattern

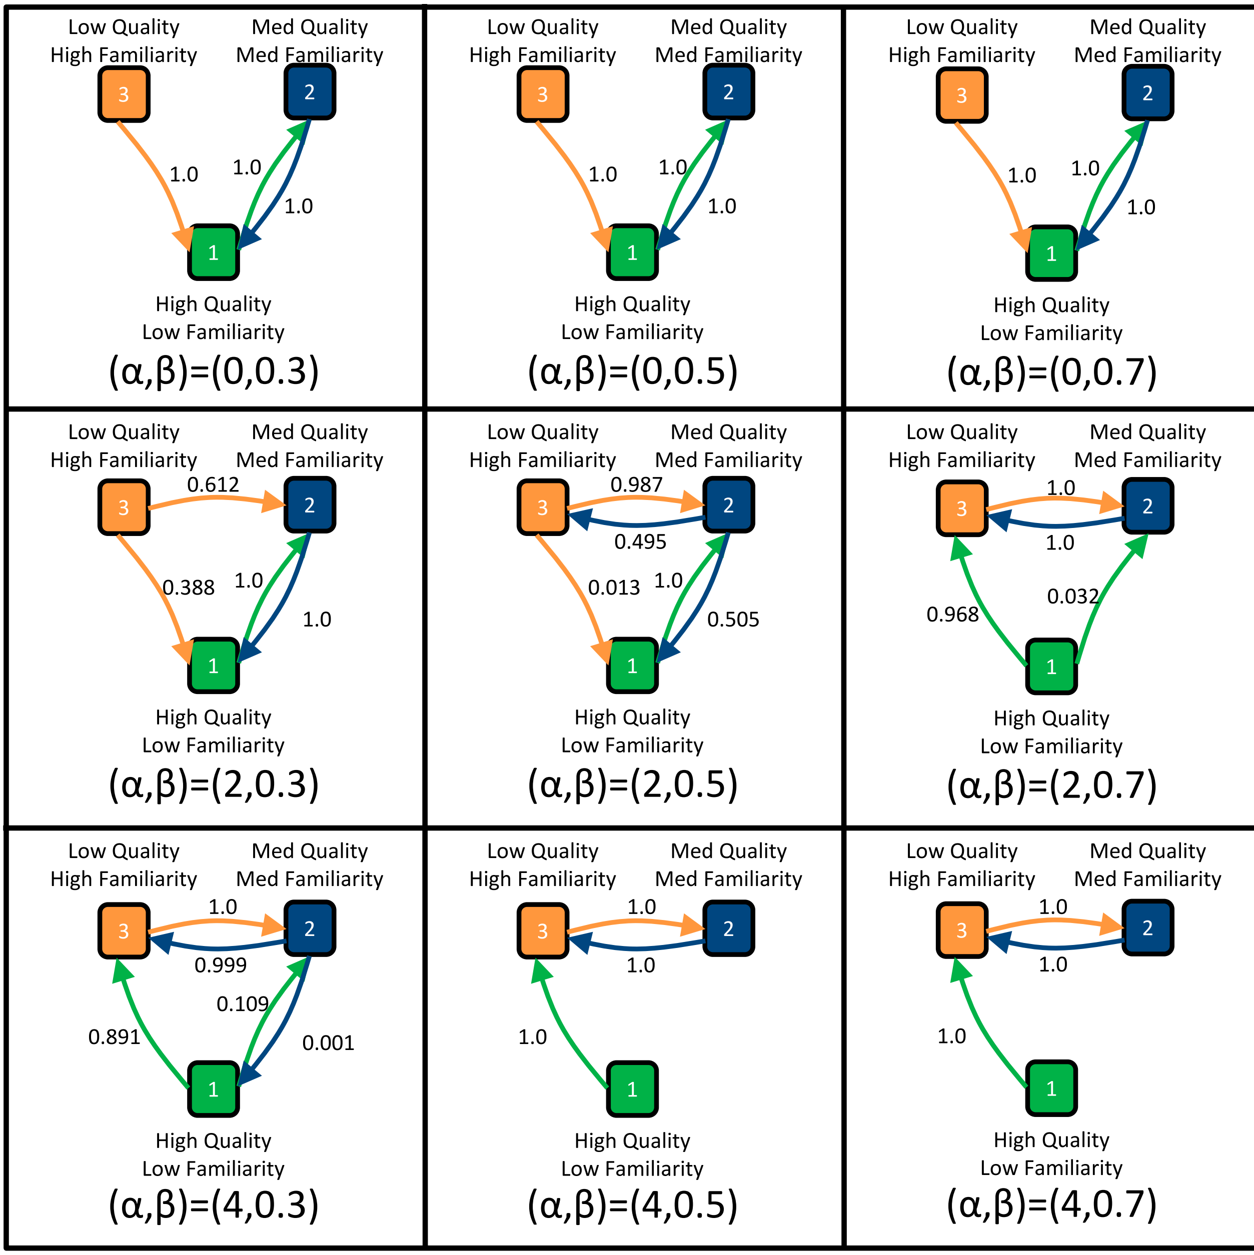

In what follows, we allow sensor qualities to be ambiguous to other sensors. We do so by assuming that any given sensor i believes that the quality of sensor $j \in \mathcal{N} \setminus \{i\}$ is contained in a set of possible values (which we refer to as the *ambiguity set*), with each possible value 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 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 previous sections, we assume that $\lfloor k \rfloor = 1$ 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 that sensor i's problem at time t is to find the targeting strategy

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

where

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

$$\text{s.t.} \quad \mathbb{P}_i^Q \left\{ \min_{\tilde{s}_{it} \in \mathbb{R}} \mathbb{E}_{\pi, S_t \sim F_{it}^{\pi}} [\tilde{s}_{it} - S_t]^2 \leq y_\epsilon \right\} \geq 1 - \epsilon. \tag{24}$$

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_{it}^{π}, where y_{it}^{π} represents the robust "cost" of a targeting strategy π. This cost is defined as the $(1-\epsilon)$-percentile (with respect to $\mathbb{P}_i^{Q}$) of sensor i's estimation squared error if it follows targeting strategy π. In (24), F_{it}^{π} 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 because of the chance-constrained optimization nature of the problem (22)–(24). This is formalized in Lemma EC.5 in Appendix H (Saghafian et al. 2018). As we will see in our numerical studies next, 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 this case, the long-run contact set, though not a singleton, can be stochastically predicted, so our earlier managerial implications about information-supplier mapping still apply. However, there are conditions under which one can restrict attention to the set of deterministic targeting policies without any loss. We present one such sufficient condition in the following proposition.

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—that is, 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 (1) quality, (2) familiarity, and (3) state can be affected by the underlying ambiguity around qualities and/or the level of optimism of sensors. We briefly summarize our observations here and refer the reader to Appendix I (Saghafian et al. 2018) 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 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 (Saghafian et al. 2018). We also observed (Study 4 in Appendix I (Saghafian et al. 2018)) 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 toward 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 online 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 Online 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 (1) order all the sensors from lowest to highest based on the variance criterion established earlier—that is, (10)—and then (2) pick the $\lfloor k \rfloor$ best-ranked—that is, lowest-variance—sensors as its targets. We prove that our key results (Propositions 1, 2, and 4 and Lemma 2) 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 is accurate in their means. The extension to a setting in which the means of these normal distributions are not correct is contained in Online 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 Online 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—that is, when state realization is simply delayed—we show that the following key results still hold: if sensor qualities are asymmetric, then (1) there exists a unique, albeit sample-path-dependent, long-run contact, (2) this contact can be determined in finite time, and (3) there exists a finite time after which the long-run communication network can be defined as a deterministic directed graph. That is, Propositions 4–6 all continue to hold; see Propositions EC.4–EC.6 in Online Appendix C for details. We then establish that these results continue to hold under a general delay with imperfect accuracy—that is, 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, so there is no state realization, except for the first period or not even in the first period.

In Appendix D (Saghafian et al. 2018), 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 the current-period estimate is analyzed in Appendix E (Saghafian et al. 2018). We establish that this nonmyopic problem lies in the general class of restless multiarmed 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 (Saghafian et al. 2018), 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 predeployment training data sets are large enough to result in sufficiently high initial precision, then network formation among sensors can be accurately studied without assuming that sensors update their own inference models after deployment. Because, in practice, firms typically use large amounts of training data predeployment, 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 predeployment, 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 because of 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 other's underlying inference models and that may not know each other's 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 kinds 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 are 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 (i.e., when i communicates with 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 also on 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 (1) the actions of the sensor-owning entities, and (2) the incentives of these entities to share information. With regard to point (1), 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 payoffs (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 point (2), 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 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 that our paper motivates further research in this area.

Endnotes

[1] Vibration is a key variable, and "vibration analysis is the most commonly used condition monitoring technique" (Syed and Pai 2016, p. 58).

[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).

[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—for example, Sparacino et al. (2007).

[5] In Appendix D (Saghafian et al. 2018), 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 extend to a setting in which state realization occurs less frequently (e.g., every $T > 1$ periods because technician inspections or data offloading is periodic because of travel or communication burdens), but target selection remains constant between realizations; this only requires a rescaling of time—that is, changing the definition of a period. Also, see Appendix G (Saghafian et al. 2018) for generalizations in which the slow-but-more-accurate estimation is not perfect and/or has a general delay of more than one period.

[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 Online Appendix B, which establishes that the key results in the paper still hold.

[9] As we will show, this implies that each sensor targets 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.

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

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

[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 (Saghafian et al. 2018).

[13] This state dependency does not arise if the underlying environment is governed by a static i.i.d. white-noise process—that is, when $\beta = 0$. In that case, it follows directly from the preceding analysis that the familiarity function $\psi_{ij,t+1}(s) = v_{ij0}^2 + f(q_i, q_j) \sum_{l=1}^{t} a_{ijt}$. 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.

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

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

[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.

References

Acemoglu D, Bimpikis K, Ozdaglar A (2014a) Dynamics of information exchange in endogenous social networks. *Theor. Econom.* 9(1):41–97.

Acemoglu D, Mostagir M, Ozdaglar A (2014b) State-dependent opinion dynamics. *2014 IEEE Internat. Conf. Acoustics Speech Signal Processing (ICASSP)* (IEEE, Piscataway, NJ), 4773–4777.

Ashton AH, Ashton RH (1985) Aggregating subjective forecasts: Some empirical results. *Management Sci.* 31(12):1499–1508.

Atanasov P, Rescober P, Stone E, Swift SA, Servan-Schreiber E, Tetlock P, Ungar L, Mellers B (2016) Distilling the wisdom of crowds: Prediction markets vs. prediction polls. *Management Sci.* 63(3):691–706.

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

Ballotta L, Schenato L, Carlone L (2019) Computation-communication trade-offs and sensor selection in real-time estimation for processing networks. Preprint, submitted November 13, https://arxiv.org/abs/1911.05859.

Bassamboo A, Cui R, Moreno A (2018) The wisdom of crowds in operations: Forecasting using prediction markets. Preprint, submitted October 27, https://dx.doi.org/10.2139/ssrn.2679663.

Bates JM, Granger CW (1969) The combination of forecasts. *J. Oper. Res. Soc.* 20(4):451–468.

BehrTech (2020) Five IoT applications for offshore monitoring in oil and gas. Accessed May 21, 2020, https://behrtech.com/blog/5-iot-applications-for-offshore-monitoring-in-oil-and-gas/.

Bren A, Saghafian S (2019) Data-driven percentile optimization for multi-class queueing systems with model ambiguity: Theory and application. *INFORMS J. Optim.* 1(4):267–287.

Chen C, Zhao XL, Li ZH, Zhu ZG, Qian SH, Flewitt AJ (2017) Current and emerging technology for continuous glucose monitoring. *Sensors* 17(1):182.

Delage E, Iancu DA (2015) Robust multistage decision making. Aleman DM, Thiele AC, eds. *The Operations Research Revolution*, INFORMS TutORials in Operations Research (INFORMS, Catonsville, MD), 20–46.

Delage E, Mannor S (2010) Percentile optimization for Markov decision processes with parameter uncertainty. *Oper. Res.* 58(1): 203–213.

Evans P, Annunziata M (2012) Industrial internet: Pushing the boundaries of minds and machines. Report, General Electric Company, Boston.

Fisher M, Raman A (1996) Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* 44(1):87–99.

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

Gaur V, Kesavan S, Raman A, Fisher ML (2007) Estimating demand uncertainty using judgmental forecasts. *Manufacturing Service Oper. Management* 9(4):480–491.

GEPower (2018) Electrical rotating machine APM overview. Accessed June 10, 2020, https://www.ge.com/news/sites/default/files/GEA33604%20Motor_Fleet_APM_Paper_0.pdf.

Hall DL, Llinas J (1997) An introduction to multisensor data fusion. *Proc. IEEE* 85(1):6–23.

Han D, Wu J, Zhang H, Shi L (2017) Optimal sensor scheduling for multiple linear dynamical systems. *Automatica J. IFAC* 75: 260–270.

International Electrotechnical Commission (IEC) (2014) Internet of things: Wireless sensor networks. White paper, International Electrotechnical Commission, Geneva, Switzerland.

IKM (2019) Condition monitoring: Oil & gas. Report, IKM, Larvik, Norway. Accessed June 10, 2020, https://www.ikm.com/getfile.php/1343855-1559300145/IKM20Selskaper/IKM20Instrutek/Nedlastninger/Datablad20og20brosjyrer/Tilstandskontroll/CM20Offshore_low.pdf.

Jernigan S, Kiron D, Ransbotham S (2016) Data sharing and analytics are driving success with IoT. *MIT Sloan Management Rev.* 58(1):1–16.

Khaleghi B, Khamis A, Karray FO, Razavi SN (2013) Multisensor data fusion: A review of the state-of-the-art. *Inform. Fusion* 14(1):28–44.

Leal Y, Garcia-Gabin W, Bondia J, Esteve E, Ricart W, Fernández-Real JM, Vehí J (2010) Real-time glucose estimation algorithm for continuous glucose monitoring using autoregressive models. *J. Diabetes Sci. Tech.* 4(2):391–403.

Lennon Olsen T, Tomlin B (2020) Industry 4.0: Opportunities and challenges for operations management. *Manufacturing Service Oper. Management* 22(1):113–122.

Liyanage LH, Shanthikumar JG (2005) A practical inventory control policy using operational statistics. *Oper. Res. Lett.* 33(4):341–348.

Marr B (2016) What everyone must know about industry 4.0. Accessed March 1, 2018, 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. Report, McKinsey Global Institute.

Mišić V, Perakis G (2021) Data analytics in operations management: A review. *Manufacturing Service Oper. Management*. Forthcoming.

Mitchell HB (2007) *Multi-Sensor Data Fusion: An Introduction* (Springer Science & Business Media, Berlin).

Momot R, Belavina E, Girotra K (2020) The use and value of social information in selective selling of exclusive products. *Management Sci.* 66(6):2610–2627.

Nemirovski A, Shapiro A (2006) Convex approximations of chance constrained programs. *SIAM J. Optim.* 17(4):969–996.

Perakis G, Roels G (2008) Regret in the newsvendor model with partial information. *Oper. Res.* 56(1):188–203.

Ran X, Chen H, Liu Z, Chen J (2017) Delivering deep learning to mobile devices via offloading. *VT/AR Network '17 Proc. Workshop Virtual Reality Augmented Reality Network* (Association for Computing Machinery, New York), 42–47.

Saghafian S (2018) Ambiguous partially observable Markov decision processes: Structural results and applications. *J. Econom. Theory* 178(11):1–35.

Saghafian S, Tomlin B (2016) The newsvendor under demand ambiguity: Combining data with moment and tail information. *Oper. Res.* 64(1):167–185.

Saghafian S, Tomlin B, Biller S (2018) The Internet of things and informmation fusion: Who talks to who? Harvard Kennedy School Working Paper 18-009, Harvard Kennedy School, Cambridge, MA. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3124417.

Sethi R, Yildiz M (2016) Communication with unknown perspectives. *Econometrica* 84(6):2029–2069.

Shi L, Zhang H (2012) Scheduling two Gauss–Markov systems: An optimal solution for remote state estimation under bandwidth constraint. *IEEE Trans. Signal Processing* 60(4):2038–2042.

Simchi-Levi D (2010) *Operations Rules: Delivering Customer Value Through Flexible Operations* (MIT Press, Cambridge, MA).

Sparacino G, Zanderigo F, Corazza S, Maran A, Facchinetti A, Cobelli C (2007) Glucose concentration can be predicted ahead in time from continuous glucose monitoring sensor time-series. *IEEE Trans. Biomed. Engrg.* 54(5):931–937.

Swaminathan JM, Tayur SR (2003) Models for supply chains in e-business. *Management Sci.* 49(10):1387–1406.

Syed S, Pai M (2016) Antifriction bearing diagnostics in a manufacturing industry: A case study. *J. Mechanical Engrg. Automation* 6(5A):58–62.

Thanagasundram S, Schlindwein FS (2006) Autoregressive order selection for rotating machinery. *Internat. J. Acoustics Vibration* 11(3):144–154.

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

Tsoukalas G, Falk BH (2020) Token-weighted crowdsourcing. *Management Sci.* 66(9):3843–3859.

Vitus MP, Zhang W, Abate A, Hu J, Tomlin CJ (2012) On efficient sensor scheduling for linear dynamical systems. *Automatica J. IFAC* 48(10):2482–2493.

Wallis KF (2011) Combining forecasts: Forty years later. *Appl. Financial Econom.* 21(1–2):33–41.

Winkler RL, Clemen RT (2004) Multiple experts vs. multiple methods: Combining correlation assessments. *Decision Anal.* 1(3):167–176.

Winnig L (2016) GE's big bet on data and analytics. *MIT Sloan Management Rev.* Reprint 57380:1–16.

Wu S, Ding K, Cheng P, Shi L (2020) Optimal scheduling of multiple sensors over lossy and bandwidth limited channels. *IEEE Trans. Control Network Systems* 7(3):1188–1200.

Xiong F, Hipszer BR, Joseph J, Kam M (2011) Improved blood glucose estimation through multi-sensor fusion. *2011 Annu. Internat. Conf. IEEE Engrg. Med. Biol. Society EMBC* (IEEE, Piscataway, NJ), 377–380.

Xu D, Li T, Li Y, Su X, Tarkoma S, Hui P (2020) A survey on edge intelligence. Preprint, submitted March 26, https://arxiv.org/abs/2003.12172.

Yang W, Chen G, Wang X, Shi L (2014) Stochastic sensor activation for distributed state estimation over a sensor network. *Automatica J. IFAC* 50(8):2070–2076.