

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

Soroush Saghafian¹, Brian Tomlin², Stephen Biller³

¹Harvard Kennedy School, Harvard University, Cambridge, MA

²Tuck School of Business at Dartmouth, Hanover, NH

³Vice President, Offering Management, Watson Internet of Things at IBM

The promised operational benefits of the Internet of Things (IoT) are predicated on the notion that better decisions will be enabled through a multitude of autonomous sensors (often deployed by different firms) providing real-time knowledge of the state of things. This knowledge will be imperfect, however, due to sensor quality limitations. A sensor can improve its estimation quality by soliciting a state estimate from other sensors operating in its general environment. Target selection (choosing from which other sensors to solicit estimates) is challenging because sensors may not know the underlying inference models or qualities of sensors deployed by other firms. This lack of trust (or familiarity) in others' inference models creates noise in the received estimate, but trust builds and noise reduces over time the more a sensor targets any given sensor. We characterize the initial and long run information sharing network for an arbitrary collection of sensors operating in an autoregressive environment. The state of the environment plays a key role in mediating quality and trust in target selection. When qualities are known and asymmetric, target selection is based on a deterministic rule that incorporates qualities, trusts, and state. Furthermore, each sensor eventually settles on a constant target set in all future periods, but this long run target set is sample path dependent and also varies by sensor. When qualities are unknown, a deterministic target selection rule may be suboptimal, and sensors may not settle on a constant target set. Moreover, the inherent targeting trade-off between quality and trust is influenced by a sensor's ambiguity attitude. Our findings shed light on the evolution of inter-firm sensor communication over time, and this is important for predicting and understanding the inter-firm connectedness and relationships that will arise as a result of the IoT.

Key words: Robust Estimation; Data Fusion; Information Sharing

History: Version: February 11, 2018

1. Introduction

It is widely believed that we are at the beginning of a new digital age. We will move from a world in which “most things have long operated dark, with their location, position, and functional state unknown or even unknowable” (Deloitte 2015, p.5) to a world in which a vast multitude of internet-connected sensors will enable near-ubiquitous monitoring of the status of assets, environments, objects, and people (Feng and Shanthikumar 2017). This Internet of Things (IoT) is expected to have profound consequences. “As almost everything—from cars to crops to conveyor belts—becomes connected, IoT is changing the way businesses operate [by generating] valuable insights to improve virtually every aspect of their operations and [enabling] innovative, new business models” IBM (2017). A report by the McKinsey Global Institute estimates that the IoT may have a “total potential economic impact of \$3.9 trillion to \$11.1 trillion per year in 2025” (McKinsey 2015, p.2),

with 63% of this attributable to operations (p.112) and with “*the real value of IoT applications coming from analyzing data from multiple sensors*” (p.104). Combining data across sensors can enable or enhance information completeness and/or estimation quality; with the latter, estimation quality, being the focus of this paper.

Information completeness refers to situations in which the states of distinct elements are combined to provide an overall state of the system. This can deliver a number of potential benefits. For example, it allows the creation of “digital twins - dynamic digital representations that enable companies to understand, predict, and optimize performance of their machines” (GE 2016, p.2). Information completeness also allows system-level optimization of “fully instrumented networks of facilities or fleets across wide geographic locations” (GE 2012, p.12). In addition to optimization across distant locations, information completeness permits granular optimization within a location. For example, Wireless Sensor Networks in agricultural fields enable micro-targeting of irrigation and pesticides by segment and potentially by plant within a field (Ling and Blexline 2017).

In this paper, we focus on the estimation-quality motive for sharing information across sensors. The IoT is only as good as the sensors that measure the states of things. Unfortunately, sensors often provide imperfect estimates due to their placement in harsh environments (Lawson 2017) or due to the inherent challenges in measuring and estimating the variable of interest. Aggregating estimates across sensors is one approach to improving estimation quality. For example, Body Sensor Networks, which are used to monitor physical activity or physiological variables, “are often characterized by error-prone sensor data, [and] the use of multi-sensor data fusion methods represents an effective solution to infer high quality information from ... noisy signals” (Gravina et al. 2017, p.68). Apple, Google, Samsung and others are reportedly developing minimally-invasive and/or non-invasive wearable bio-sensors to measure blood glucose (BG) levels; technologies include tear-based sensors, salivary-based sensors, and skin-penetrating optical sensors. Due to accuracy and precision concerns, it is envisioned that these technologies will require “frequent calibration against direct BG data” obtained through invasive means (Chen et al. 2017, p.8) and that outputs from multiple sensors may be combined to generate an improved BG estimate (Xiong et al. 2011). More generally, sensor fusion is concerned with how best to efficiently transfer and aggregate signals across a pre-determined set of sensors that share information. The concept of information “fusion across sensors [that] nominally measure the same property [to] reduce or eliminate noise and errors” (Mitchell 2007, p.4-5) is envisioned as a remedy for sensor quality concerns in a wide range of IoT settings including energy, environmental monitoring, infrastructure, and various other industrial

applications (IEC 2015).

Our paper explores a different but vital question in the era of IoT-enabled information fusion: which sensors should “talk” to which sensors? That is, from which sensors should any given sensor solicit information? This is an essential question for business in light of the vision of the IoT as a vast collection of sensors deployed by many firms, with sensor-enabled devices “making simple decisions on their own and becoming as autonomous as possible” (Marr 2016). The International Electrotechnical Commission reports that “complex systems, for instance city-wide sensor networks, will not necessarily be owned by one group. Also, more often than not, more than one organization will operate in the same system” (IEC 2015, p.33). Jernigan et al. (2016) report that “two-thirds of the respondents to [their] survey who are actively working on IoT projects collect data from and/or send data to the customers, suppliers or competitors” (p.4). As the number of sensors (in any given environment) grows with increased IoT penetration, one might anticipate that estimation quality concerns will be remedied through frequent information aggregation across a very large number of sensors. However, there are important factors that complicate that vision.

A sensor might not be able to solicit estimates (so as to improve its own estimate) at any given point in time from all other sensors in its environment due to cost and technical considerations. From a cost perspective, firms understand that their sensor-derived information is valuable to other firms, and therefore they can charge other firms for access to their sensor estimates.¹ Technical constraints can arise due to communication channel capacity and bandwidth limitations as well as individual sensor energy consumption concerns.² As such, even if a sensor is in an environment with a large number of other sensors, it may be limited in the number of sensors from which it can solicit estimates at any given instant. Faced with such a constraint, intuition might suggest that a sensor should target the higher quality sensors, where quality refers to the precision of sensor measurements.

As we will show in this paper, that intuition breaks down in a world where sensors choose targets from among sensors deployed by other firms. A sensor generates an estimate of the (time-varying) variable of interest (hereafter referred to as “state”) by combining its current observation with an inference model of how the state evolves randomly over time.³ An implication of autonomous

¹ For example, Terbine is developing a platform that “will act as a broker mediating between those generating data and those wanting to consume it” (Gibbs 2016) so as to enable firms to “monetize the IoT/sensor data that [they] are already collecting” (<http://terbine.com/what-is-terbine.html>).

² “The wireless modem that serves as the link between a wireless sensor and the outside world is the largest consumer of energy on the sensor. Therefore, it is important to be selective about how the wireless communication channel is used”, (Swartz et al. 2010, p.3).

³ The sensor literature often uses the term “process model” as opposed to “inference model”, but we use the latter

inter-firm sensor communication is that a sensor might not know the inference model used by another sensor (typically deployed and trained by another firm)⁴: it will have to interpret the other sensor's estimate based on its own beliefs about the other sensor's inference model. Trust evolves over time the more a particular sensor is targeted, and past communication therefore influences future targeting. Another very important implication of inter-firm sensor targeting is that a sensor might not have full information about the quality of other sensors.

As firms increasingly rely on sensor-connected devices to make autonomous decisions, it is vital that management have the ability to predict how their devices will interact with devices deployed by other firms. What kind of communication network—who talks to who—will develop over time? Understanding the pattern of sensor-driven communication across firms is extremely important because such communication “deepens existing relationships between organizations and forges new relationships,” (Jernigan et al. 2016, p.4). The challenge of dynamically determining which subset of sensors to target so as to improve one's own estimate of the state of a dynamic random environment is the fundamental question addressed in this paper. We consider a collection of autonomous sensors operating in a common environment that 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. Higher quality sensors are less noisy: they can generate more precise signals of the state. A sensor knows its own quality but may not know the quality of other sensors. Furthermore, each sensor has its own private inference model based on its pre-deployment training and tuning and its own understanding (i.e., trust) of the other sensors' inference models.

In each period, after observing its own private noisy signal of the state of the environment, each sensor chooses a subset of other sensors to target (i.e., from which sensors to solicit state-of-the-environment estimates) so as to generate an improved state estimate. The updating of the sensor's state estimate depends on (a) its understanding of the targeted-sensor's inference model, which update over time the more that sensor is targeted, and (b) the qualities of both the targeted and the targeting sensors. We start our analysis by assuming that sensors know the qualities of all other sensors. We then relax the known-quality assumption using a robust optimization approach known as percentile optimization. This enables the sensors to make robust target selection decisions while being ambiguous about the qualities of the other sensors. For expositional clarity, we focus on the

throughout this paper as it is a more common term in both the management sciences and the estimation theory literatures.

⁴ Training is typically done independently by each deploying firm with their own independent data sets and algorithms. This results in sensors deployed by different firms having different inference models.

setting in which a sensor can choose only one target during each period because this is the setting in which choosing the right target is the most important. However, all of our results extend easily to a setting in which multiple simultaneous targets are allowed.

Among other results, we establish that the state of the environment plays a key role in determining the weights placed on quality versus trust when selecting a target in any given period. Furthermore, because targeting builds future trust, the current state also influences future target selection. We establish that when qualities are known and asymmetric, each sensor will eventually target a single sensor in all future periods but this long run target can vary by sensor. State dependency means that these long run targets are sample path dependent, and hence, even for each particular sensor, the long run target can vary depending on the realization of state over time. Nevertheless, we show that the long run communication network that forms between sensors can be fully defined at time zero as a random directed graph. Interestingly, we also prove that this random directed graph can be fully characterized as a deterministic directed graph after a finite time. That is, when sensor qualities are known and asymmetric, one can deterministically characterize the long run communication targets of all sensors after observing state evolution for a finite time.

When qualities are not common knowledge (i.e., sensors face ambiguity with respect to other sensor qualities), randomizing across some subset of sensor may be optimal in the long run even along a given sample path. Nevertheless, we provide an intuitive sufficient condition under which a deterministic targeting policy is optimal. We also establish that a sensor's ambiguity attitude (which will depend on the deploying firm) plays an important role in target selection.

The rest of the paper is organized as follows. The most relevant literature is discussed in §2. The base model is described in Section §3. Analysis and results are presented in §4 and §5. The extension to unknown sensor quality is developed and analyzed in §6. Conclusions are discussed in §7. All proofs are provided in the appendix.

2. Literature

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

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. (2015) empirically explores the effect of group size on forecast accuracy, finding that aggregation across larger groups improves accuracy. The notion that aggregation of a large number of estimates can improve estimation quality—nowadays 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, Palm and Zellner 1992, Winkler and Clemen 2004, Wallis 2011, Acemoglu et al. 2014, Atanasov et al. 2016). Through that lens, one can view our work as exploring a related but different question: when each individual in the crowd wants to improve his or her own estimate (but cannot ask everyone in the crowd) then who in the crowd should an individual target?

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

Multi-sensor data fusion, defined by Mitchell (2007)[p.3.] as “the theory, techniques and tools which are used for combining sensor data, or data derived from sensory data, into a common representational format ... so that it is, in some sense, better than would be possible if the data sources were used individually”, emerged as a problem domain in the 1990s due to the U.S. military’s desire to enable more-complete or higher quality surveillance of geographic areas. It has

since grown to encompass diverse applications in artificial intelligence, robotics, and environmental, equipment and health monitoring (Hall and Llinas 1997). Sensor fusion is typically focused on developing efficient and effective architectures and data processing techniques (Hall and Llinas 1997, Mitchell 2007, Khaleghi et al. 2013, Hall and Llinas 1997) for sharing across a given network of sensors, whereas our paper, motivated by the autonomous nature of sensors in the IoT, explores the question of which subset of sensors any given sensor should target.

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), Chu et al. (2008), Perakis and Roels (2008), Ramamurthy et al. (2012), Saghafian and Tomlin (2016), and references therein. For some general theoretical results on the percentile optimization approach that we utilize, we refer interested readers to Nemirovski and Shapiro (2006), Delage and Mannor (2010), and references therein.

3. The Model

We model a setting in which a collection $\mathcal{N} \triangleq \{1, 2, \dots, N\}$ of autonomous sensors operates in a common environment. In what follows, we describe the environment, individual sensor measurement and state estimation, sensor collaboration, and finally the target selection problem whereby each sensor chooses from which other sensors to solicit state estimates.

Environment: The environment is defined by a state variable $S \in \mathbb{R}$ whose discrete-time evolution is governed by a first-order autoregressive process (AR(1)):

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

for $t = 1, 2, \dots, \infty$, where $\tilde{\epsilon}_t$ i.i.d. normal white-noise random variables with mean 0 and variance normalized to 1. We adopt an AR(1) model as the most parsimonious one that allows the environment to exhibit autoregressive behavior, a common feature in many situations.

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

$$\Gamma_{it} = S_t + \epsilon_{it}, \quad (2)$$

where ϵ_{it} are i.i.d. normal white noises with mean 0 and variance $1/(q_i)^2$, with q_i representing the quality of sensor i . 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 deploying firms’s training algorithms and data sets prior to deployment—which is given by

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

where $\hat{\alpha}_i$ and $\hat{\beta}_i$ are sensor i ’s estimates of the process parameters α and β . At the very end of each period, the true state becomes publicly known (i.e., is revealed to each sensor), and the system moves to the next period.⁵ Our analysis and findings readily extend to a setting in which state realization occurs less frequently (e.g., every $T > 1$ periods) but target selection remains constant between realizations. This only requires a re-scaling of time (i.e., changing the definition of period).

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

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

and

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

respectively, where S_{it} represents sensor i ’s signal-updated belief (distribution) about the current state S_t . The higher the quality of sensor i , the more weight it places on its signal when updating its mean belief, and the larger the associated variance reduction.

Information Sharing and Sensor Collaboration: Each sensor $i \in \mathcal{N}$ is aware of all the other sensors in the environment. All sensors in the collection \mathcal{N} are willing to collaborate in the following

⁵ For example, as discussed in §1, it is thought that non-invasive blood-glucose sensors will require periodic inputs from precise but invasive readings.

manner: in each period t , after all sensors have formed updated beliefs based on their private signals (according to (4) and (5) above), any sensor $j \in \mathcal{N}$ is willing to share its best estimate of state which is its updated mean predication of state $E[S_{jt} | \Gamma_{jt} = \gamma_{jt}]$ with any other sensor i that requests it.⁶ 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 different firms typically train their sensors (pre-deployment) differently using different algorithms and data sets. 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.⁷ In this setting, parameters $v_{ij0} > 0$ and $w_{ij0} > 0$ represent the initial *trust* (or familiarly) that sensor i has in sensor j 's inference model. A setting in which sensor i fully knows sensor j 's inference model parameters can be obtained by setting $v_{ij0} = w_{ij0} = \infty$. To gain insights, in our base model we assume that sensor qualities q_i are common knowledge to all $i \in \mathcal{N}$, but this is relaxed in §6.

Target Selection: In each period t , after updating its state estimate based on its private signal as in (4) and (5) above, each sensor i chooses a set of sensors from which to request state estimates (i.e., their updated mean beliefs about the state). We do not model the actions of devices associated with sensors but implicitly assume that the action payoff is increasing in the quality of the state estimate. In choosing which sensors to target, sensor i selects those sensors that will most improve its own estimate. By most improvement, we mean that sensor i 's resulting updated state distribution gives the lowest expected squared error of estimation.⁸ In particular, sensor i solves the following optimization problem in each period t :

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

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. We define $k = b/c$, and refer to it as

⁶ Our work easily extends to a setting in which j will only collaborate with some subset of \mathcal{N} .

⁷ The extension of our analyses to a setting in which the mean of these normal distributions is not correct is relatively straightforward.

⁸ As we will show, this implies that each sensor targets those sensors that provide it with the most information about the state. All our results hold for any targeting objective function that is strictly increasing in the expected squared error of estimation.

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 during each period, i.e., $\lfloor k \rfloor = 1$, because this is when choosing the right target is most important. Our results readily extend to the case of $\lfloor k \rfloor > 1$ as we will discuss in the next section.

4. Preliminaries

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

We begin by establishing that the target selection problem in (6) above is equivalent to one in which sensor i selects as its target the sensor that provides i with the most informative signal about the current state, where a less noisy signal (i.e., one with a lower variance) is more informative.⁹ 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 trust in sensor j 's inference model. In particular, given its privately generated signal Γ_{jt} in period t , sensor j provides sensor i with its best current estimate of state, which is $\mathbb{E}[S_{jt}|\Gamma_{jt}]$, i.e., its updated/latest expected belief about the current state S_t . Now, from sensor i 's perspective, $\mathbb{E}[S_{jt}|\Gamma_{jt}]$ is formed according to:

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

which is similar to (4) above 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 $\hat{\beta}_j$. Because $\Gamma_{jt} = S_t + \epsilon_{jt}$ from (2), this value $\mathbb{E}[S_{ijt}|\Gamma_{jt}]$ provides sensor i with the following noisy signal regarding the state S_t :

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

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

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

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

⁹ Note that the information entropy of any normally distributed random variable depends only on its variance.

$\Xi_{ijt}(s_{t-1}) = \hat{\alpha}_{ijt} + \hat{\beta}_{ijt}s_{t-1}$, where its dependence on the prior state value s_{t-1} is explicitly noted. It has a variance of

$$\text{Var}[\Xi_{ijt}(s_{t-1})] = \frac{1}{v_{ijt}^2} + \frac{s_{t-1}^2}{w_{ijt}^2}, \quad (10)$$

where $1/v_{ijt}^2$ and $1/w_{ijt}^2$ are the variances at time t of sensor i 's beliefs about sensor j inference model parameters $\hat{\alpha}_j$ and $\hat{\beta}_j$. Defining the precision $\psi_{ijt}(s_{t-1}) \triangleq 1/\text{Var}[\Xi_{ijt}(s_{t-1})]$, it follows from (9) and (10) that

$$\begin{aligned} \sigma_t^2(i, j, s_{t-1}) &= \frac{1}{q_j^2} + \frac{1}{q_j^4} \left(\frac{1}{v_{ijt}^2} + \frac{s_{t-1}^2}{w_{ijt}^2} \right) \\ &= \frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4}. \end{aligned} \quad (11)$$

Under a variance reduction objective, sensor i chooses as its target in period t :

$$\begin{aligned} j_{it}^* &\triangleq \arg \min_{j \in \mathcal{N} \setminus \{i\}} \sigma_t^2(i, j, s_{t-1}) \\ &= \arg \min_{j \in \mathcal{N} \setminus \{i\}} \left\{ \frac{q_j^2 + 1/\psi_{ijt}(s_{t-1})}{q_j^4} \right\}. \end{aligned} \quad (12)$$

The following result establishes that the original target selection problem in (6) is equivalent to the variance reduction target selection (12); that is, both objectives result in the same target.¹⁰

PROPOSITION 1 (Target Selection and Variance Reduction). *If channel capacity $[k] = 1$, then under (6), $a_{ijt}^* = \mathbb{1}_{\{j=j_{it}^*\}}$, where j_{it}^* is given by (12), 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 $[k]$ other sensors that provide the lowest variance of signal to firm i . That is, it chooses the $[k]$ most-informative sensors (from its perspective) and solicits their state estimates.

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 *trust function* that sensor i has for sensor j at time t , and we refer to $\psi_{ijt}(s_{t-1})$ as the *trust value*, i.e., the trust function evaluated at the prior state $s = s_{t-1}$. We use the term *trust* to convey the notion that higher values imply less errors in sensor i 's understanding of sensor j 's underlying inference model. Importantly, as we will establish below, sensor i does not need to separately update its beliefs over time about parameters $\hat{\alpha}_{ijt}$ and $\hat{\beta}_{ijt}$ of sensor j 's inference model; it suffices to update the trust function $\psi_{ijt}(s)$.

To operationalize the target selection problem (12), we now examine how any given sensor's trust function with respect to some other sensor evolves over time. In particular, we develop the mechanism through which the time- t trust function is updated to that at time $t + 1$, i.e., how

¹⁰ Without loss of generality, we assume ties in (6) and (12) are broken by choosing the sensor with the lower index.

$\psi_{ijt}(s)$ updates to $\psi_{ij,t+1}(s)$. To that end, we first note that it follows from (10) that the initial trust function is given by

$$\psi_{ij1}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2}. \quad (13)$$

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

$$(1 + q_j^2) \mathbb{E}[S_{jt} | \Gamma_{jt}] = \Xi_{ijt}(s_{t-1}) + q_j^2 (S_t + \epsilon_{jt}). \quad (14)$$

There are two independent sources of noise in this signal: (a) the noise in sensor i 's own estimate of the current state S_t , which has variance of $1/(1 + q_i^2)$ (see (5)), and (b) the inherent white noise ϵ_{jt} in sensor j measurement, which has a variance of $1/q_j^2$. Thus, based on (14), the variance in the signal's noise is given by $\text{Var}[q_j^2 (S_t + \epsilon_{jt})] = q_j^4 / (1 + q_i^2) + q_j^2$. Using Bayesian updating, we have the following result.

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

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

$$f(q_i, q_j) \triangleq \frac{(1 + q_i^2)}{q_j^2 (1 + q_i^2 + q_j^2)}. \quad (15)$$

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

$$\psi_{ij,t+1}(s) = \frac{v_{ij0}^2 w_{ij0}^2}{w_{ij0}^2 + v_{ij0}^2 s^2} + f(q_i, q_j) \sum_{l=1}^t a_{ijl}. \quad (16)$$

Intuitively, sensor i 's trust function for sensor j changes from time t to time $t + 1$ if, and only if, i targets j at time t , i.e., $a_{ijt} = 1$. Moreover, if i targets j then the gain in i 's trust in j does not depend on the state: the gain is given by $f(q_i, q_j)$, which we refer to as the *stickiness factor*. It is noteworthy, however, that the gain depends on both the sender's (j 's) and the receiver's (i 's) qualities. More importantly, it follows from (16) that to calculate the current trust 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.

5. Communication Networks: Who Targets Who?

With the equivalent target selection problem and trust 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 trust value $\psi_{ijt}(s_{t-1})$ for each sensor j ; see the targeting criterion (12). The attractiveness of j as a potential target for i depends on the trust value $\psi_{ijt}(s_{t-1})$. The trust 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.¹¹

5.1. Initial Target Selection

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

To answer these questions, let $r \triangleq q_l/q_h$ denote the quality ratio of sensors l and h . By definition, $0 < r \leq 1$. Using (12) and (13), it follows that sensor i strictly prefers targeting the lower quality sensor (l) if, and only if,

$$\frac{(r q_h)^2 + \left(\frac{1}{v_{il0}}\right)^2 + \left(\frac{s_0}{w_{il0}}\right)^2}{(r q_h)^4} < \frac{q_h^2 + \left(\frac{1}{v_{ih0}}\right)^2 + \left(\frac{s_0}{w_{ih0}}\right)^2}{q_h^4}, \quad (17)$$

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 $\hat{\beta}_j$. From (17), it can be seen that i strictly prefers to target l if, and only if,

$$c_0 + c_1 s_0^2 > 0, \quad (18)$$

where

$$c_0 = (r q_h)^2 (r^2 - 1) + \left(\frac{r^2}{v_{ih0}}\right)^2 - \left(\frac{1}{v_{il0}}\right)^2, \quad (19)$$

and

$$c_1 = \left(\frac{r^2}{w_{ih0}}\right)^2 - \left(\frac{1}{w_{il0}}\right)^2. \quad (20)$$

We note that c_0 reflects a tension between the difference in sensor qualities and the differences in i 's noise in its beliefs about 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 noise in its beliefs about the inference model parameters $\hat{\beta}_h$ and $\hat{\beta}_l$. As (18) shows, state matters in initial target selection through this c_1 term. The following result presents the conditions under which sensor i strictly prefers to sacrifice quality for trust. By an appropriate swapping of labels l and h , it can also

¹¹ This state dependency does not arise if the underlying environment is governed by a static i.i.d. white noise process, i.e., when $\beta = 0$. In that case, it follows directly from the above analysis that the trust function $\psi_{ij,t+1}(s) = v_{ijt}^2 = v_{ij0}^2 + f(q_i, q_j) \sum_{l=1}^t a_{ijl}$. 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.

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:*

- (i) $c_0 \leq 0$, $c_1 > 0$, and $|s_0| > \sqrt{-c_0/c_1}$,
- (ii) $c_0 > 0$ and $c_1 < 0$, and $|s_0| \leq \sqrt{-c_0/c_1}$, or
- (iii) $c_0 > 0$ and $c_1 \geq 0$.

This proposition highlights the interconnected roles that (a) sensor qualities, (b) trusts, and (c) state play in target selection. Intuitively, if sensor i has more trust in its beliefs about the high-quality sensor's inference model parameters (i.e., $v_{ih0} \geq v_{il0}$ and $w_{ih0} \geq w_{il0}$), then the high-quality sensor is the inherently more attractive target regardless of state. This is reflected in the above proposition by the fact that $c_0 < 0$ and $c_1 < 0$ in this case and, therefore, sensor h is preferred. On the other hand, if sensor i has more trust in its beliefs about 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. This tradeoff between quality and trust depends on the state (parts (i) and (ii) of Proposition 3) unless the trust advantage of the lower quality sensor compared to the higher quality one is so large that it makes the lower quality sensor the preferred target regardless of the state (part (iii) of Proposition 3).

As the quality ratio r increases from 0 to 1 (all else held constant) there are at most three distinct regions of target selection, as illustrated in Figure 1.¹² When the quality ratio r is low, i.e., sensor h is of much higher quality than l , then h dominates l , i.e., h is targeted in all states. This h -dominating region always exists, but it does not cover the entire range $0 < r \leq 1$ unless $v_{ih0} \geq v_{il0}$ and $w_{ih0} \geq w_{il0}$, i.e., the high quality sensor is more trusted for both parameters. In contrast, when the quality ratio is high, i.e., sensor qualities relatively similar, then l dominates h , i.e., l is targeted in all states. This l -dominating region exists if, and only if, $v_{ih0} < v_{il0}$ and $w_{ih0} < w_{il0}$, i.e., the low quality sensor is more trusted for both parameters. Importantly, there is an intermediate range of the quality ratio r (that extends to $r = 1$ if the trust ranking differs across v and w) in which state matters and the indifference curve $|s_0| = \sqrt{-c_0/c_1}$ completely characterizes target selection. Figure 1 illustrates an instance with parameters for which Proposition 3(i) applies. In this case, a high absolute value of state induces sensor i to emphasize trust over quality such that it targets sensor l .

¹² A complete closed-form analytical characterization of the region thresholds exists but it is algebraically cumbersome and not included for reasons of space.

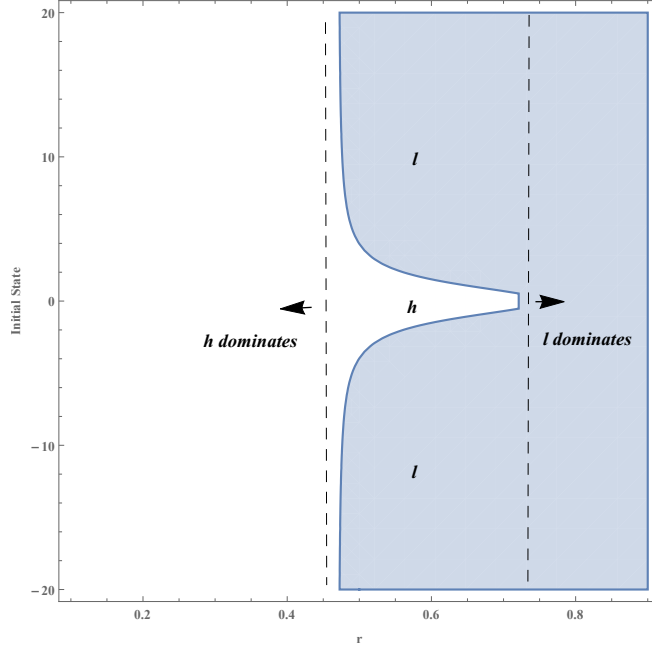


Figure 1 Initial selection between a higher quality sensor (h) and a lower quality sensor (l). [Proposition 3 (i) applies in the intermediate region.]

In contrast, when the absolute value of state is low, quality matters more than trust, and i targets h . The reverse holds if case (ii) applies. When this intermediate region exits, then case (i) (i.e., high state favors high trust sensor) applies over this entire intermediate region if $r > \sqrt{w_{ih0}/w_{il0}}$, but case (ii) (i.e., high state favors high quality sensor) applies over this entire intermediate region otherwise.

5.2. Target Evolution and Long Run Target Selection

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

The stickiness factor $f(q_i, q_j)$ determines the gain in trust that results each time i targets j . It

is strictly increasing in q_i and strictly decreasing in q_j ; see (15). This has two implications worth noting. First, all else equal, if two sensors i_h and i_l with $q_{i_h} > q_{i_l}$ both target some other sensor j then the resulting gain in trust in j is higher for i_h than for i_l . In other words, higher quality sensors can build trust in other sensors more rapidly than can lower quality sensors. Second, all else equal, if two sensors j_h and j_l with $q_{j_h} > q_{j_l}$ are potential targets for some other sensor i , then i 's potential gain in trust is lower for i_h than for i_l . Put differently, lower quality sensors result in larger trust 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 Trusts that Vary by Sensor): Consider the case in which any given sensor i has a common initial trust in all other sensors $j \in \mathcal{N} \setminus \{i\}$, i.e., $v_{ij0} = v_{i0}$ and $w_{ij0} = w_{i0}$ for all j . This common initial trust 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 Trusted 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 trusts in both of $\hat{j}(i)$'s inference model parameters than in any other sensor's parameters. It follows from part (iii) of Proposition 3 that $\hat{j}(i)$ dominates all other sensors (from the perspective of i) at time 1. Because dominance is preserved (Lemma 1), the communication network at any time (including the long run) is the same across all sample paths: regardless of state realizations, each sensor i always targets its initially most-trusted sensor $\hat{j}(i)$.

In general, however, sensors may differ in their qualities and any given sensor may have heterogeneous trusts in 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 $\mathcal{S}_\infty \triangleq \{s_0, s_1, s_2, \dots\}$ denote a long run sample path, i.e., the realization of states as time approaches infinity. Similarly, we denote by $\mathcal{S}_t \triangleq \{s_0, s_1, s_2, \dots, s_t\}$ a sample path up to time t . We also let $\mathcal{S}_\infty'^t$ denote a sample path that is equivalent to \mathcal{S}_∞ up to time t , but one which may deviate from \mathcal{S}_∞ afterwards: $\mathcal{S}_\infty'^t \triangleq \mathcal{S}_t \cup \{s'_{t+1}, s'_{t+2}, \dots\}$. To examine the long run networks that may arise, we first introduce the following definition.

DEFINITION 2 (Long Run Trustees). Given a sample path \mathcal{S}_∞ , the set of long run trustees of sensor i is:

$$\mathcal{T}_i(\mathcal{S}_\infty) \triangleq \{j \in \mathcal{N} \setminus \{i\} : \lim_{t \rightarrow \infty} \psi_{ijt}(s_{t-1}) = \infty | s_0, s_1, s_2, \dots \in \mathcal{S}_\infty\}. \quad (21)$$

REMARK 1 (Infinitely-Often Communication). It immediately follows from (16) that, along any sample path \mathcal{S}_∞ , 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 trust of some sensor i in these equal-quality sensors, there might exist some sample paths along which the long run set of trustees of sensor i includes more than one sensor and sensor i keeps alternating between the sensors in its long run set of trustees 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}_∞ : (a) $\mathcal{T}_i(\mathcal{S}_\infty)$ is a singleton, i.e., $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$, and (b) the unique long run trustee in $\mathcal{T}_i(\mathcal{S}_\infty)$ can be identified in the almost sure sense in finite time, i.e., $\mathcal{T}_i(\mathcal{S}_\infty) = \mathcal{T}_i(\mathcal{S}_\infty'^{t^*})$ 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*, i.e., a directed graph with given probabilities assigned to each link ij that indicate

the probability that j will be the long run target for i . Furthermore, there exists a finite time after which the graph can be defined as a deterministic directed graph, i.e., with all probabilities being zero or one, that fully specifies the long run target for each sensor.

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

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

The above lemma states that a sensor j' will not be targeted by a sensor i if (a) there is another sensor j of a higher quality than j' , and (b) sensor i 's trust in 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 trust level in sensor j' . Thus, the above lemma holds regardless of how much i trusts j' at time t : if i 's trust in j passes the fixed threshold, then j' will not be targeted by i . This in turn allows us to show that, when sensor qualities are asymmetric (defined below), the set of long run trustees of each sensor i along any sample path \mathcal{S}_∞ 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 Trustee**). *If sensor qualities are asymmetric, then given any sample path \mathcal{S}_∞ , $|\mathcal{T}_i(\mathcal{S}_\infty)| = 1$ for all $i \in \mathcal{N}$.*

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

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

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

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

The above results allow us to characterize the long run network of communications.¹³ 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

$$p_{ij} \triangleq \sum_{\forall \mathcal{S}_\infty: j \in \mathcal{T}_i(\mathcal{S}_\infty)} Pr(\mathcal{S}_\infty). \quad (22)$$

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

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

As noted earlier, all of the above results extend readily to a setting in which a sensor can simultaneously target $\lfloor k \rfloor > 1$ other sensors. The long run trustee set of each sensor will, however, contain $\min\{\lfloor k \rfloor, |\mathcal{N}| - 1\}$ sensors.

5.3. Numeric Studies of Long Run Network

All else equal, it follows from the above analysis that the at time $t = 0$ the probability that some sensor j will be the long run trustee of some other sensor i increases in the quality of sensor j and in i 's initial trust in j 's inference model parameters; that is, p_{ij} , defined in (22), increases in q_j , v_{ij0} and w_{ij0} . It is less clear how the underlying environment influences long run target selection. We now present the results of two numerical studies that examine, respectively, the roles that the underlying state dynamics and the initial state distribution play in long run communication network formation. In both studies, for ease of illustration, we consider a collection of three sensors, i.e., $\mathcal{N} = \{1, 2, 3\}$. The following sensor characteristics are adopted for each study. Sensor qualities

¹³ The extension of the above results to settings with non-asymmetric sensors is straightforward, although as noted earlier, Proposition 4 may no longer hold.

are asymmetric and decreasing in sensor labels, with $\underline{q} \triangleq (q_j : j \in \mathcal{N}) = (9.8, 9.6., 9.4)$. The initial trusts v_{ij0} and w_{ij0} are specified by the following symmetric matrix:

$$[v_{ij0}]_{i,j \in \mathcal{N}} = [w_{ij0}]_{i,j \in \mathcal{N}} = \begin{pmatrix} \infty & 1.6 & 2.4 \\ 0.8 & \infty & 2.4 \\ 0.8 & 1.6 & \infty \end{pmatrix}. \quad (23)$$

In other words, any two different sensors i and k have the same initial trusts about the third sensor j . This matrix implies that initial trusts are increasing in sensor labels. Combined with the above quality-labeling scheme, it then follows that sensor 1 is the most attractive from a quality perspective but that sensor 3 is the most attractive from an initial trust perspective. Sensor 2 lies in between sensors 1 and 3 in that it represents mid values of both quality and trust. By design, these parameters ensure that no sensor is initially dominant from the perspective of any other sensor. Also, these parameters ensure that target selection in the initial period is given by Proposition 3 (i); that is, the intermediate region of Figure 1 applies. For later reference, it follows that sensor 1 initially targets the lower quality but higher-trust sensor (3 from 1's perspective) if, and only if, the initial state $|s_0| > 4.32$. Sensor 2 initially targets the lower quality but higher-trust sensor (3 from 2's perspective) if, and only if, the initial state $|s_0| > 2.27$. Sensor 3 initially targets the lower quality but higher-trust sensor (2 from 3's perspective) if, and only if, the initial state $|s_0| > 1.60$.

For each problem instance in each study, the time-0 random directed graph $\vec{G}(\mathcal{N}, \mathcal{E}, \mathcal{P})$ was generated by simulating 1,000 sample paths, each representing a distinct realization of the underlying AR(1) process over time. That is, at time $t = 0$, the probability that j will be the eventual long run target of i (see p_{ij} defined in (22)) is estimated using the outcomes of the 1,000 sample paths.¹⁴ The initial state s_0 of the underlying AR(1) process is also randomly drawn from a distribution specified in each study below.

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

Before exploring the effect of α and β , we first discuss the long run communication network for the middle-middle case of $(\alpha, \beta) = (2, 0.5)$. This is depicted in Figure 2, where the probability (at time $t = 0$) that j will be the eventual long run target of i , i.e., p_{ij} , is shown on the link (i, j) .

¹⁴ The number of simulated sample paths were chosen so that the point estimations for p_{ij} values have a low enough standard error, and hence, provide reliable estimation.

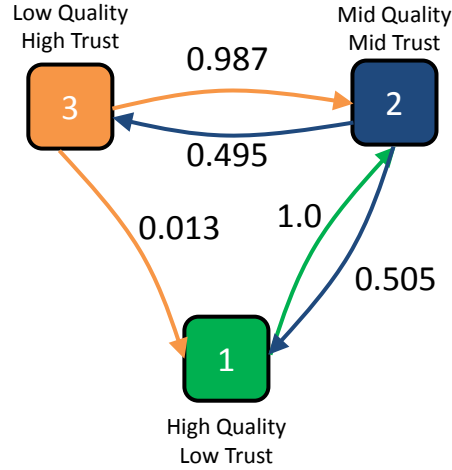


Figure 2 The long run communication network in Study 1 for $(\alpha, \beta) = (2, 0.5)$.

Consider sensor 3’s targeting choice. As noted earlier, sensor 3 initially targets 2 (3’s higher-trusted lower quality sensor) if, and only if, the initial state satisfies $|s_0| > 1.60$. This has a probability of 0.981. Furthermore, sensor 3’s trust gain in 2 (if 2 targeted) is higher than its trust gain in 1 (if 1 targeted). Therefore, 3’s preference for 2 typically grows over time such that 2 becomes the long run trustee across almost most but not all sample paths: observe that $p_{32} = 0.987$. Trust does not always win out over quality, because sample paths can occur for which sensor 1 is targeted frequently enough (relative to 2) to render it eventually dominant.

Let us next consider sensor 2’s targeting choice. It chooses between the two extreme sensors: sensor 1 (the highest quality but lowest trust sensor) and sensor 3 (the lowest quality highest trust sensor). The quality advantage of 1 and the trust advantage of 3 are somewhat balanced such that each sensor has a reasonable likelihood of becoming sensor 2’s long run trustee; observe that $p_{23} = 0.495$. This is in spite of the fact that the initial state distribution very heavily favors sensor 3. (Sensor 2 initially targets 3 if, and only if, the initial state $|s_0| > 2.27$, which has a probability of 0.93.) Finally, let us consider sensor 1’s targeting choice. Observe that $p_{12} = 1$ in all cases. Although the higher quality lower-trust sensor (2 in 1’s case) is not initially dominant, it eventually becomes the long run trustee across all sample paths. The trust advantage that 3 has is not sufficient to overcome the quality advantage of 2 frequently enough to ever make 3 dominant.

Figure 3 presents the long run communication networks for all nine cases of (α, β) . We observe a strong effect in both α and β . When α and β are both low, the long run trustee of each sensor is always the high-quality sensor, i.e., 1 eventually always targets 2, and 2 and 3 eventually always target 1.¹⁵ When α and β are both high, the long run trustee of each sensor is always the high-trust

¹⁵ We note that at higher values of β , e.g., $\beta = 0.9$ (not shown), a lower quality sensor has a small probability of

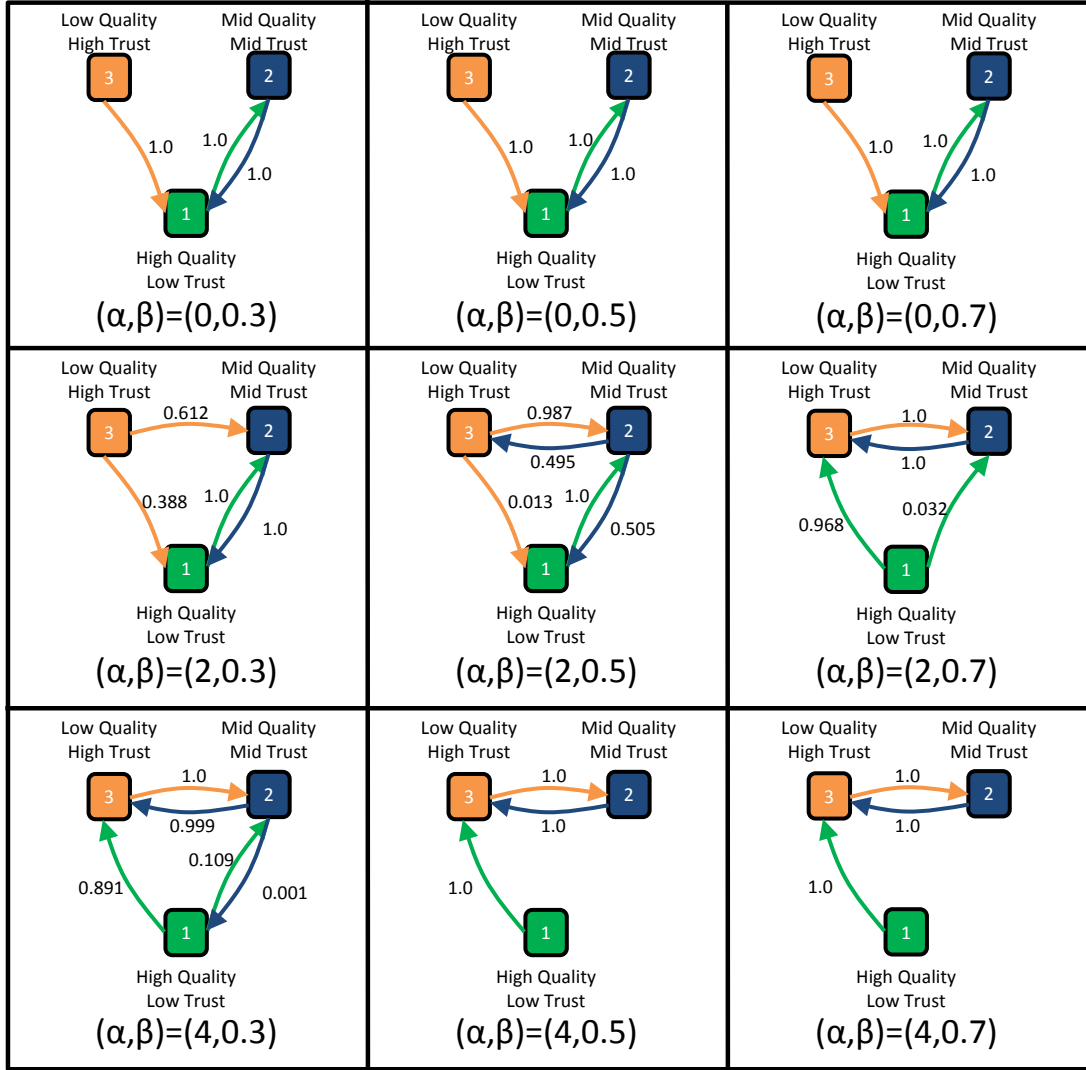


Figure 3 The role of state dynamics.

sensor, i.e., 1 and 2 target 3, and 3 targets 2. At intermediate values of α and β , the long run trustee of each sensor typically depends on the sample path: for some sample paths, the higher quality sensor wins and for others the higher-trust sensor wins. However, we observe that the winner is more likely to be the higher-trusted sensor as either α and β increases. The reason for this (α, β) effect is twofold. One, high state values favor higher trusted targets because Proposition 3 (i) applies (by study design). Two, higher state values are more likely to occur on any given sample path as either α or β increases. Therefore, higher-trust sensors are increasingly favored as α and/or β increases. Interestingly, we also note that α typically has a stronger effect than β .

Study 2 (The Effect of Initial State): In this study, we fix the environment's AR(1) parameters as $(\alpha, \beta) = (2, 0.5)$ and consider three different $N(\mu, \sigma)$ distributions for the initial state: (a) $(\mu, \sigma) = (2, 2)$; (b) $(\mu, \sigma) = (10, 2)$, and (c) $(\mu, \sigma) = (18, 2)$. The long run communication (targeting) network being the long run trustee even when $\alpha = 0$.

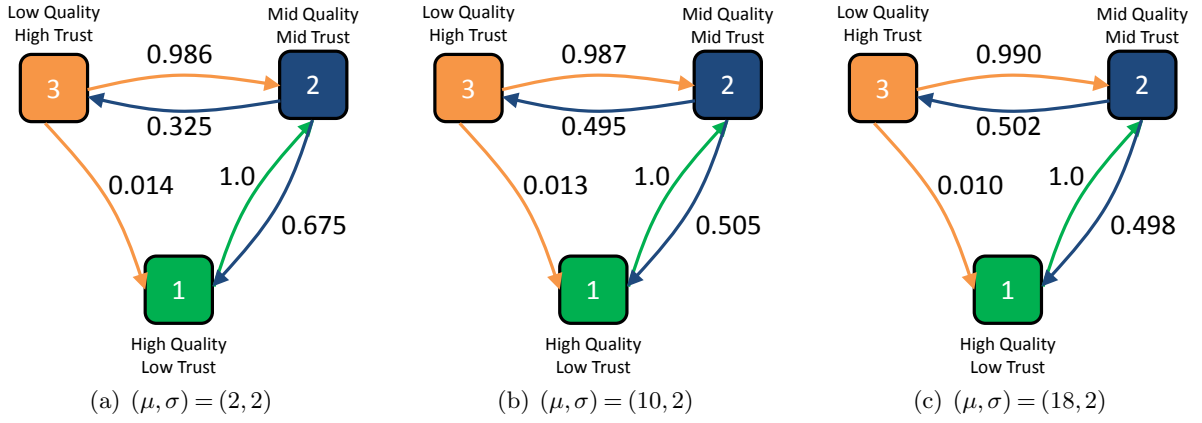


Figure 4 The role of initial state in long run communication network. S_0 has a $N(\mu, \sigma)$ distribution.

for each case is presented in Figure 4. Regardless of the initial state's realization, the long run distribution of the environment has mean of 4.0 and a standard deviation of 1.15 because $(\alpha, \beta) = (2, 0.5)$.¹⁶

Let us first consider sensor 3's targeting choice. It initially targets 2 (3's higher-trust but lower-quality sensor) if, and only if, $|s_0| > 1.60$. This has a probability of 0.58 in case (a), 0.99999 in (b), and 1 in (c). Therefore, as we move from case (a) to (b) to (c) the starting state distribution increasingly favors initial selection of 2 (the more trusted sensor), and this slightly increases 2's probability of being the long run trustee. Let us next consider sensor 2's targeting choice. It initially targets 3 (2's higher-trusted but lower quality choice) if, and only if, $|s_0| > 2.27$. This has a probability of 0.45 in case (a), 0.9999 in (b), and 1 in (c). Therefore, as we observed for Sensor 3, as we move from case (a) to (b) to (c) the starting state distribution increasingly favors initial selection of the more-trusted sensor and this increases the probability of that sensor being the long run trustee. Finally, observe that for sensor 1 the long run trustee is 2 (1's higher quality target) with probability 1 (i.e., regardless of the sample path realization) in all cases. Although, as with the other sensors, there is a higher probability of the more-trusted sensor being initially targeted as we move from case (a) to (b) to (c), this has no effect on the long run trustee probability because quality eventually wins over initial trust in all sample paths.

In summary, the initial state can impact the long run trustee probability through its impact on initial selection and the resulting trust gain. That is, due to the stickiness factor, initial state can create a momentum that might last for ever. However, in comparing Study 1 and Study 2, we observe that the effect of the initial state is less strong than the effect of the underlying state

¹⁶ For (α, β) the long run distribution of the state is normal with mean $\alpha/(1-\beta)$ and standard deviation of $1/\sqrt{1-\beta^2}$ because the process is scaled so that the random zero-mean noise terms have a variance of one.

dynamics, i.e., the (α, β) parameters. This is because the (α, β) parameters influence the state realization in every period, whereas the initial state's impact on future states diminishes over time.

6. Sensors of Unknown Qualities

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

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

$$\pi_{it}^* = \arg \inf_{\pi \in \Pi_i} y_{it}^\pi \quad (24)$$

where

$$y_{it}^\pi = \inf_{y_\epsilon \in \mathbb{R}_+} y_\epsilon \quad (25)$$

s.t.

$$\mathbb{P}_i^Q \left\{ \min_{\tilde{s}_{it} \in \mathbb{R}} \mathbb{E}_{\pi, S_t \sim F_{it}^\pi} [\tilde{s}_{it} - S_t]^2 \leq y_\epsilon \right\} \geq 1 - \epsilon. \quad (26)$$

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 the sensor i 's estimation squared error if it follows targeting strategy π . In (26), F_{it}^π is the posterior distribution of sensor i 's belief about the state after applying the targeting strategy π at time t .¹⁷ The $\epsilon \in [0, 1]$ parameter represents the level of *optimism*, where

¹⁷ 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 trust function at time t which in turn depends on q_j values), but these dependencies are suppressed for ease of notation.

$\epsilon = 0$ 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 (24) need not be deterministic in general: a randomized strategy might outperform any deterministic strategy due to the *chance constrained* optimization (a.k.a percentile optimization) nature of problem (24)-(26). As we will see, this randomization may cause a sensor to have a long run set of trustees 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 trust with more than one other sensor in the long run, and go back and forth between them infinitely often (along any given sample path).

To observe that a randomized policy can be better than any deterministic one, note that under a deterministic policy that prescribes targeting j (almost surely) we can write (26) as:

$$Pr\{V_j \leq y_\epsilon\} \geq 1 - \epsilon \quad (27)$$

where $V_j \triangleq \Sigma^2(i, j, s_{t-1}, Q_j)$ is a random variable with realization $\sigma^2(i, j, s_{t-1})$ (defined in (11)), and Q_j is a random variable with realization q_j .¹⁸ That is, sensor i solves a robust counter part to the original variance reduction problem (see, e.g., (12)) in which instead of connecting to the sensor j that has the minimum $\sigma^2(i, j, s_{t-1})$ it connects to the sensor j that has the minimum

$$F_{V_j}^{-1}(1 - \epsilon), \quad (28)$$

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

$$F_\Xi^{-1}(y) \triangleq \inf\{z \in \mathcal{Z} : F_\Xi(z) \geq y\}. \quad (29)$$

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

$$y_d^*(\epsilon) = \min_{j \in \mathcal{N} \setminus \{i\}} F_{V_j}^{-1}(1 - \epsilon), \quad (30)$$

and sensor i connects to sensor

$$j_{it,d}^* = \arg \min_{j \in \mathcal{N} \setminus \{i\}} F_{V_j}^{-1}(1 - \epsilon). \quad (31)$$

¹⁸ Note that due to the fact that Q_j is a random variable for sensor i , the trust value of i to j is also a random variable. Thus, the trust value $\psi_{i,j,t+1}(s_t)$ has a realization which can be calculated based on (16) for each $Q_j = q_j$ that belongs to the ambiguity set considered by sensor i .

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

$$y_r(\epsilon) = F_{\bar{V}}^{-1}(1 - \epsilon). \quad (32)$$

Since, unlike $F_{V_j}^{-1}$, $F_{\bar{V}}^{-1}$ depends on the joint distribution of all Q_j 's defined by \mathbb{P}_i^Q (for all sensor j 's where $j \in \mathcal{N} \setminus \{i\}$), it can be the case that $y_r(\epsilon) < y_d^*(\epsilon)$. That is, one can find a feasible randomized targeting strategy for sensor i that is strictly better than any deterministic policy. Because the randomized policy that resulted in $y_r(\epsilon)$ is not necessarily optimal, it follows that the optimal policy π_{it}^* defined in (24) may not be deterministic.

However, there are conditions under which one can restrict attention to the set of deterministic policies without any loss. Below, we provide one such sufficient condition.

PROPOSITION 7 (Deterministic Communication). *Suppose that at period t for all $j \in \mathcal{N} \setminus \{i\}$ we have $V_{\tilde{j}} \leq_{s.t.} V_j$ for some $\tilde{j} \in \mathcal{N} \setminus \{i\}$. Then π_{it}^* defined in (24) prescribes that sensor i targets sensor \tilde{j} at period t almost surely regardless of ϵ .*

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

Next, we numerically explore how the tradeoffs we observed in our earlier results and experiments between (a) quality, (b) trust, and (c) state can be affected by the underlying ambiguity around qualities and/or by the level of optimism of sensors.¹⁹ In both of the following studies (Studies 3 and 4 below), we assume that a collection of three sensors operates in an environment where the underlying AR(1) model has parameters $(\alpha, \beta) = (2, 0.7)$. We assume that sensor $i \neq j$ believes that sensor j 's quality lies in the range $(0, 2q_j)$ with all values in this range equally likely and

¹⁹ We restrict our attention to the set of deterministic policies to gain clear insights, and avoid extra levels of complexities that are caused by randomized policies.

where $\underline{q} \triangleq (q_j : j \in \mathcal{N}) = (9.8, 9.6., 9.4)$. Thus, the average perceived quality and ambiguity range both decrease in the sensor label, such that Sensor 1 (3) has both the highest (lowest) maximum possible quality and the highest (lowest) average perceived quality. That is, while a higher-label sensor is perceived to have a higher quality on average, a lower-label sensor's quality is known with less ambiguity. In what follows, the known-quality counterpart problem sets the qualities at the exact q_j values given by $\underline{q} \triangleq (q_j : j \in \mathcal{N}) = (9.8, 9.6., 9.4)$, which is (a) similar to values used in our earlier studies, and (b) corresponds to the average perceived qualities in the unknown-quality case.

Study 3 (Common Initial Trusts) In this study, we assume initial trust values are common among sensors (set at $v_{ij0} = w_{ij0} = 2$), and compare the cases of known and unknown qualities. For the case of known qualities, we analytically established earlier that target selection is deterministic (i.e., sample path independent) and time-invariant in the case of common initial trusts (see earlier Special Case 1): the highest-quality sensor targets the second-highest quality sensor, all other sensors target the highest quality sensor. This is illustrated in Figure 5(a) in which the link value is 1 on $1 \rightarrow 2$; $2 \rightarrow 1$; and $3 \rightarrow 1$. The long run trustee may not be deterministic, however, when qualities are unknown: see Figure 5(b) which presents the long run communication network for an optimism parameter of $\epsilon = 0.2$.

Observe from the link values that in the long run each sensor typically (but not always) targets the lowest-label sensor available to it, indicating that each sensor's long run target is more likely to be the one for which the ambiguity is lowest even though such a sensor's average perceived quality is the lowest. However, this preference towards lower ambiguity (which comes at the cost of targeting a lower average quality sensor) is violated on some sample paths. A value of $\epsilon = 0.2$ implies that sensors are quite pessimistic (i.e., ambiguity-averse) in target selection, and hence they put a lot of weight on avoiding targeting a potentially low-quality sensor. Given our construction of the ambiguity sets, lower-label sensors have lower likelihoods of low quality values, and that is why sensors tend to target low-label sensors. However, because there is some chance that the lower label sensor may be the one with the lower quality (e.g., higher labeled sensors have a higher maximum possible quality), we observe that this long run low-label selection does not occur along all sample paths.

Study 4 (The Effect of the Optimism Level): We now explore the impact of the optimism parameter ϵ on long run target selection. We consider a setting where initial trusts vary across sensors. In particular, we use the same initial trust matrix as used in our earlier known-quality studies (see (23)). Therefore, the study replicates the known-quality instance in Figure 3(f) except

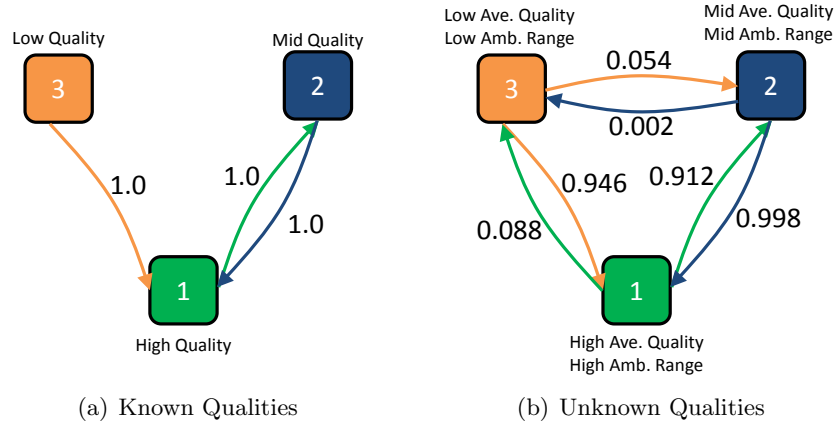


Figure 5 With unknown qualities, the communication network is stochastic even with identical initial trusts.

that there is now quality ambiguity; sensors can't assume that qualities are exactly at their means as was the case with known qualities. Figure 6(a) replicates the known-quality communication network, i.e., Figure 3(f), for ease of comparison. Figures 6(b) and 6(c) present the unknown-quality communication networks for optimism parameters $\epsilon = 0.2$ and $\epsilon = 0.9$, respectively. One immediate observation is that the link probabilities differ when qualities are unknown. This is because target selection now must consider ambiguities in addition to qualities and trusts. Moreover, we observe that the optimism parameter strongly influences long run target selection. Sensors are more pessimistic (optimistic) about potential sensor qualities when ϵ is low (high) and this in turn influences the emphasis placed on trust 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. Therefore, organizational attitudes towards ambiguity will impact target selection and the resulting communication network that evolves over time.

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

7. Conclusions

Much of the promise of the IoT stems from the idea that better operational decisions will be made because a vast array of sensors will enable almost real-time knowledge of the state of things. This knowledge will often be imperfect due to inherent precision limitations of sensors. Information fusion, in particular, the sharing of estimates across sensors, can help improve estimation quality. However, sensors cannot necessarily solicit information from all other sensors due to cost and technical considerations; they may be limited in their number of targets at any given instant.

We characterize the initial and long run communication network—who talks to who—for an

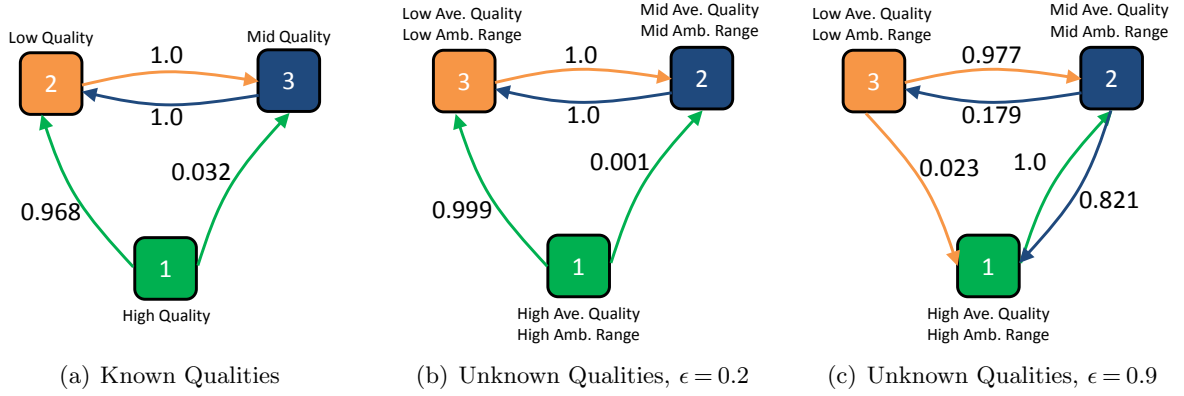


Figure 6 Communication network differs when qualities unknown and depends on the optimism parameter ϵ

arbitrary collection of sensors that do not know each others' underlying inference models and that may not know each others' qualities. We establish that the state of the environment plays a key role in determining the weights placed on quality and trust (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 as well as realization of states over time (sample path). These long run targets are sample path dependent, because state value influences the weight sensors put on trust versus quality. We prove that the random directed graph that characterizes the long run stochastic communication network at time $t = 0$ becomes deterministic (i.e., with links that have 0 or 1 probabilities) after a finite time.

When qualities are not common knowledge, we show that a firm's attitude towards ambiguity (through its choice of its sensor's optimism parameter) can play an important role in target selection. Also, we establish that different to the case of known (and asymmetric) qualities, a sensor might not eventually settle on one particular target even along a given sample path: randomizing across some subset in every period (including long run) may be optimal.

Our work sheds light on what kind of communication networks develop over time, and this enables managers to make predictions about how their devices will interact with devices from other firms. Understanding the evolution of inter-firm sensor communication is a very important aspect of understanding the inter-firm connectedness and relationships that will arise as a result of the IoT. This specific research could be extended in a number of directions. For example, 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. Also, the environment might evolve according to a more general model than the AR (1) model we used to generate insights. More

broadly, the IoT presents many opportunities to explore how to improve and exploit information quality and information completeness in operations and supply chain settings. Finally, developing an understanding of communications among sensors deployed by different firms may also help inform how to structure contracts and relationships (informal and formal alliances) between firms.

References

- Acemoglu, Daron, Kostas Bimpikis, Asuman Ozdaglar. 2014. Dynamics of information exchange in endogenous social networks. *Theoretical Economics* **9**(1) 41–97.
- Ashton, Alison Hubbard, Robert H Ashton. 1985. Aggregating subjective forecasts: Some empirical results. *Management Science* **31**(12) 1499–1508.
- Atanasov, Pavel, Phillip Rescober, Eric Stone, Samuel A Swift, Emile Servan-Schreiber, Philip Tetlock, Lyle Ungar, Barbara Mellers. 2016. Distilling the wisdom of crowds: Prediction markets vs. prediction polls. *Management Science* **63**(3) 691–706.
- Bassamboo, Achal, Ruomeng Cui, Antonio Moreno. 2015. The wisdom of crowds in operations: Forecasting using prediction markets. *Working Paper, Available at SSRN: <https://ssrn.com/abstract=2679663>* .
- Bates, John M, Clive WJ Granger. 1969. The combination of forecasts. *OR* **20**(4) 451–468.
- Chen, Cheng, Xue-Ling Zhao, Zhan-Hong Li, Zhi-Gang Zhu, Shao-Hong Qian, Andrew J Flewitt. 2017. Current and emerging technology for continuous glucose monitoring. *Sensors* **17**(1) 182.
- Chu, Leon Yang, J George Shanthikumar, Zuo-Jun Max Shen. 2008. Solving operational statistics via a bayesian analysis. *Operations Research Letters* **36**(1) 110–116.
- Delage, Erick, Shie Mannor. 2010. Percentile optimization for markov decision processes with parameter uncertainty. *Operations research* **58**(1) 203–213.
- Deloitte. 2015. Inside the internet of things: A primer on the technologies building the IoT .
- Feng, Qi, George Shanthikumar. 2017. How research in production and operations management may evolve in the era of big data. *Production and Operations Management (Forthcoming)* .
- Fisher, Marshall, Ananth Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research* **44**(1) 87–99.
- Gaur, Vishal, Saravanan Kesavan, Ananth Raman, Marshall L Fisher. 2007. Estimating demand uncertainty using judgmental forecasts. *Manufacturing & Service Operations Management* **9**(4) 480–491.
- GE. 2012. Industrial internet: Pushing the boundaries of minds and machines .
- GE. 2016. The digital twin: Compressing time-to-value for digital industrial companies .

- Gibbs, Mark. 2016. Terbine tests marketplace for IoT data. <https://www.networkworld.com/article/3046173/big-data-business-intelligence/terbine-tests-marketplace-for-iot-data.html> Accessed January 4 2018.
- Gravina, Raffaele, Parastoo Alinia, Hassan Ghasemzadeh, Giancarlo Fortino. 2017. Multi-sensor fusion in body sensor networks: State-of-the-art and research challenges. *Information Fusion* **35** 68–80.
- Hall, David L, James Llinas. 1997. An introduction to multisensor data fusion. *Proceedings of the IEEE* **85**(1) 6–23.
- IBM. 2017. What is the IoT? <https://www.ibm.com/internet-of-things/learn/what-is-iot/> Accessed December 18 2017.
- IEC. 2015. *Internet of Things: Wireless Sensor Networks*. International Electrotechnical Commission.
- Jernigan, Stephanie, David Kiron, Sam Ransbotham. 2016. Data sharing and analytics are driving success with IoT. *MIT Sloan Management Review* **58**(1).
- Khaleghi, Bahador, Alaa Khamis, Fakhreddine O Karray, Saiedeh N Razavi. 2013. Multisensor data fusion: A review of the state-of-the-art. *Information Fusion* **14**(1) 28–44.
- Lawson, S. 2017. Look before you leap: 4 hard truths about IoT. *IT World, March 23, 2017* .
- Ling, G., B. Blextine. 2017. Precision farming increases crop yields. *Scientific American, June 26, 2017* .
- Liyanage, Liwan H, J George Shanthikumar. 2005. A practical inventory control policy using operational statistics. *Operations Research Letters* **33**(4) 341–348.
- Marr, B. 2016. What everyone must know about industry 4.0. *Forbes, June 20, 2016* .
- McKinsey. 2015. *The Internet of Things: Mapping the value beyond the hype*. McKinsey Global Institute.
- Mitchell, Harvey B. 2007. *Multi-sensor data fusion: An introduction*. Springer Science & Business Media.
- Nemirovski, Arkadi, Alexander Shapiro. 2006. Convex approximations of chance constrained programs. *SIAM Journal on Optimization* **17**(4) 969–996.
- Palm, Franz C, Arnold Zellner. 1992. To combine or not to combine? Issues of combining forecasts. *Journal of Forecasting* **11**(8) 687–701.
- Perakis, Georgia, Guillaume Roels. 2008. Regret in the newsvendor model with partial information. *Operations Research* **56**(1) 188–203.
- Ramamurthy, Vivek, J George Shanthikumar, Zuo-Jun Max Shen. 2012. Inventory policy with parametric demand: Operational statistics, linear correction, and regression. *Production and Operations Management* **21**(2) 291–308.
- Saghafian, Soroush, Brian Tomlin. 2016. The newsvendor under demand ambiguity: Combining data with moment and tail information. *Operations Research* **64**(1) 167–185.

- Sethi, Rajiv, Muhamet Yildiz. 2016. Communication with unknown perspectives. *Econometrica* **84**(6) 2029–2069.
- Simchi-Levi, David. 2010. *Operations Rules: Delivering customer value through flexible operations*. MIT Press.
- Swaminathan, Jayashankar M, Sridhar R Tayur. 2003. Models for supply chains in e-business. *Management Science* **49**(10) 1387–1406.
- Swartz, R Andrew, Jerome P Lynch, Stephan Zerbst, Bert Sweetman, Raimund Rolfes. 2010. Structural monitoring of wind turbines using wireless sensor networks. *Smart Structures and Systems* **6**(3) 183–196.
- Wallis, Kenneth F. 2011. Combining forecasts—forty years later. *Applied Financial Economics* **21**(1-2) 33–41.
- Winkler, Robert L, Robert T Clemen. 2004. Multiple experts vs. multiple methods: Combining correlation assessments. *Decision Analysis* **1**(3) 167–176.
- Xiong, Feiyu, Brian R Hipszer, Jeffrey Joseph, Moshe Kam. 2011. Improved blood glucose estimation through multi-sensor fusion. *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*. IEEE, 377–380.