Service Networks with Open Routing and Procedurally Rational Customers

November 2023

Forthcoming in *Production and Operations Management*

Andrew E. Frazelle
Jindal School of Management, The University of Texas at Dallas, Richardson, TX 75080
andrew.frazelle@utdallas.edu

Tingliang Huang
Haslam College of Business, The University of Tennessee, Knoxville, Knoxville, TN 37996
thuang7@utk.edu

Yehua Wei
The Fuqua School of Business, Duke University, Durham, NC 27708
yehua.wei@duke.edu

Self-interested customers' form of reasoning and its consequences for system performance affect the planning decisions of service providers. We study *procedurally rational* customers—customers who make decisions based on a sample containing anecdotes of the system times experienced by other customers. Specifically, we consider procedurally rational customers in two-station service networks with open routing, i.e., customers can choose the order in which to visit the stations. Because some actions may be less represented in the population, a given customer may not succeed in obtaining anecdotes about all possible actions. We introduce a novel sampling framework that extends the procedurally rational framework to incorporate the possibility that a customer may not receive any anecdotes for one of the actions; in this case, the customer uses a prior point estimate in lieu of the missing anecdotes. Under this framework, we study the procedurally rational equilibrium in open routing. We show first that as the sample size grows large, customers' estimates become more accurate, and the procedurally rational equilibrium converges to the fully rational equilibrium (which is also socially optimal). We then uncover two main findings. First, we obtain bounds on the distance between the procedurally rational and fully rational equilibrium, aiding operational planning and showing the rate of convergence to the fully rational outcome as the sample size of anecdotes of each individual customer grows. Second, if customers obtain anecdotes of both actions with high probability, then the equilibrium will approximate the fully rational outcome, despite the sampling error inherent to procedural rationality.

Key words: open routing, bounded rationality, anecdotal reasoning, behavioral operations, queueing

1. Introduction

In service systems with multiple stations, customers are often free to choose their routes through the network, creating an *open-routing* environment. In particular, in environments such as theme parks, shopping malls, and catered receptions, there are few structural restrictions on the sequence of stations that customers visit, enabling strategic routing choices. Practical case studies of service systems that fit the criteria for open routing, i.e., having multiple stations which need not be visited in a fixed sequence, are performed in Baron et al. (2016) and Shtrichman et al. (2001). Baron et al. (2016) study a medical clinic where multiple tests must be performed but the order is mostly irrelevant, and Shtrichman et al. (2001) discuss an army recruitment office where the recruits must submit to multiple independent evaluations. In these works the routing is flexible but centralized, so customers cannot self-select their routes. Systems with both open routing and self-interested customers have been studied by Parlaktürk and Kumar (2004) in a queueing model under steady state, and Arlotto et al. (2019) in a model where customers are present before the start of the service. Additionally, Honnappa and Jain (2015) study the "network concert queueing game," which involves customers choosing their arrival times to a queueing network as well as their routes through the network.

The existing literature on strategic open routing has assumed customers to be fully rational. However, an open-routing service network is a complex system, and even to compute (much less implement) the fully rational equilibrium requires customers to know all system parameters and then map out the interaction of multiple queues over all possible collective routing decisions. Even if some customers are sophisticated, the question remains what each customer assumes about the rationality of the others. Therefore, it is perhaps likely that customers will not behave like fully rational agents in such systems. Instead, they may adopt simple heuristics to decide on their routes, exhibiting bounded rationality. Bounded rationality has been studied in several strategic queueing contexts in recent years (see, e.g., Huang et al. 2013 and other references in Section 2), mostly in queueing systems with only one station. However, we are not aware of any work that studies bounded rationality in a multi-station, open-routing service network, which we believe to be the type of nuanced setting in which customers are likely to exhibit such behavior.

The present work aims to address this gap. Relaxing the assumption of full rationality opens up a world of possible modeling choices to incorporate bounded rationality. Some common choices in the behavioral operations literature include quantal response equilibrium (Su 2008), reference dependence (Wu et al. 2015, Yang et al. 2018), and procedural rationality (Ren et al. 2018). Modeling bounded rationality is complex in almost any context. In this paper, we propose a tractable model of bounded rationality that extends the framework of procedural rationality. The notion of procedural rationality was introduced by Osborne and Rubinstein (1998) as an alternative to the traditional Nash equilibrium concept. Their $S(K)$-reasoning framework (K representing a sample size) is meant to more closely resemble human behavior by eschewing strong assumptions about players' sophisticated reasoning and even their awareness of the game's parameters, instead allowing customers to "sample" possible moves and learn from the outcomes. In addition to Osborne and Rubinstein (1998), Spiegler (2006) also motivates his model of anecdotal reasoning from (Tversky and Kahneman 1971) that reported experiments in which human decision-makers "over-infer from small samples." The so-called "sample naivete" phenomenon is also documented and studied in the psychology literature; see, e.g., Juslin et al. (2007) and references therein. Finally, in an operational context, Tong and Feiler (2017) adapt these concepts from psychology to study the "naive intuitive statistician," which is closely related to procedural rationality but does not involve reasoning about others' actions. In their model, a planner draws a "mental sample" of finite size from a probability distribution and mistakenly treats the properties of that sample as those of the true distribution.

The framework of procedural rationality has the benefit of requiring only one parameter: the number of "anecdotes" that each player takes for a given action. Moreover, it emulates a familiar formula from daily life of reasoning via anecdotes: asking colleagues to recommend a doctor after moving to a new city; asking a friend about her experience at a newly opened restaurant; or sampling different routes on a daily commute. For this reason, procedural rationality is also called *anecdotal reasoning* because individuals reason based on small samples or anecdotes, which constitute "word of mouth." Aside from simplicity and its resemblance to human behavior, another important reason for choosing the procedural rationality framework is that the base model naturally extends to incorporate several important features in this paper's setting. In our open-routing model featuring two stations, the anecdotes are categorized based on the routes through the service network, and customers can receive anecdotes for each of the two different routes to assess and compare. However, it is natural to expect that anecdotes for a certain route that is less represented in the population may be harder to sample than others, a feature that is not captured by existing models of procedural rationality.

To further clarify this concept, let us consider the following real-world example. A guest planning her trip to Disney World might post on social media, asking about others' experiences with different routes through the park. One such post from July 2021, in the Facebook Group "Walt Disney World Tips and Tricks"[1] (a group with over 600,000 members as of August 2022) reads "We're headed to WDW. . . our first day we will be at Magic Kingdom. . . We plan to [arrive at] *rope drop*, what route should we take at the park. . . ?"[2] Note that "rope drop" refers to the park's opening time when a significant number of guests enter simultaneously and make their routing decisions. This scenario closely resembles our model's setting, where customers are present at the start of service. To further illustrate the concept within the context of our two-station model, let us imagine a hypothetical guest who makes a similar post, in a simplified system with only two rides (say, A and B): "which ride did you visit first, A or B, and how long were your wait times at both rides?" The replies to this post can be assumed to depend on the population proportion of routes, and if one of the routes is rare in the population, then it may not appear at all in the comments.

In the procedural rationality framework, and motivated by practical examples of customer reasoning about routing decisions like the one above, we seek a model with four important facets, namely one that: (i) is parsimonious (ideally, one or few parameters); (ii) captures the fact that the probability

[1] https://www.facebook.com/groups/disneytipsandtricks (accessed 08/01/2022)

[2] https://www.facebook.com/groups/disneytipsandtricks/posts/1930298137135681 (accessed 04/05/2022)

of obtaining both types of anecdotes is significantly affected by the population proportion; (iii) is reasonably aligned with customer behavior; and (iv) explicitly models the decision process for a customer who does not receive any anecdotes for one of the routes.

For the open-routing network, we focus on a model with two stations in which customers must visit both stations but can freely choose the sequence of service. Customers want to minimize the total amount of time that they spend in the system, but they exhibit procedural rationality in their reasoning about wait times. In our model, the customer decision-making process resembles the literature on procedural rationality (see, e.g., Osborne and Rubinstein 1998, Spiegler 2006), where each customer obtains anecdotes from others and bases their decision on these anecdotes. However, in this literature, the availability to customers of one or more anecdotes about each alternative is usually taken for granted. By contrast, a crucial contribution of our work to the procedural rationality literature is that we explicitly model the process customers use to sample the anecdotes, including the dependence of this process on the prevalence of each route in the population. This also entails handling the case in which a customer obtains all anecdotes from only one route and thus must make a decision *without any anecdote* from the other route. Consequently, our model necessitates that customers possess a *prior estimate* regarding the expected system time for each route.

We consider two related sampling processes for customers to obtain anecdotes: random route anecdotes and general route anecdotes. The random route anecdotes process mirrors the uniform random sampling from the entire population, where each customer obtains a sample with a fixed total number of anecdotes. The number of anecdotes for a specific route is a binomial random variable, with the total number of anecdotes as the trial parameter and the fraction of the population choosing that route as the probability parameter. We also consider general route anecdotes, where the numbers of anecdotes from the two routes follow a general joint distribution. Under either sampling process, it is possible for a customer's sample to contain anecdotes from only one of the two routes, in which case she has no anecdotes from the other route. To handle this case, a customer is endowed with a prior estimate of the system time of each route, which she uses if she does not have anecdotes for the route.

For tractability, we study a fluid model with a continuum of infinitesimal customers. Fluid models have been commonly used in both rational queueing games (Akan et al. 2012) and the procedural rationality framework (Spiegler 2006). In addition, our numerical study in Section 6 demonstrates that the insights gleaned from our fluid model indeed translate to the discrete setting. We note that while Arlotto et al. (2019) studied the fully rational counterpart of our model with discrete customers, the analysis of procedurally rational customers presents difficulties because the customers reason based on anecdotes. It is much easier for the customers than inferring the true expected system time, but it adds randomness to their *reasoning* as well as their decisions, making the discrete customer setting exceedingly cumbersome to analyze. Importantly, this randomness is fundamentally different from that in a mixed-strategy Nash equilibrium, in which players' actions are random but their reasoning about expected utility is deterministic and perfectly accurate. By contrast, procedurally rational customers have noise in their estimates of the system time of each route, creating randomness in their reasoning itself. In the fluid model, the customer behavior is still complex, but since the aggregate behavior evolves deterministically, a careful analysis allows us to characterize the equilibria.

Our study makes a theoretical contribution to the procedural rationality literature by incorporating randomness in the number of each type of anecdote and by explicitly treating the case where a customer has no anecdotes for a particular alternative. We characterize the response function of procedurally rational customers and subsequently examine the procedurally rational equilibrium. Specifically, we compare the procedurally rational equilibrium with results from the fully rational model. In the latter, Arlotto et al. (2019) observed that customers *herd* by all choosing the same route through the network via a pure strategy. We uncover two primary findings, which we detail next.

First, we derive a closed-form bound on the difference between the procedurally rational equilibrium and herding for random route anecdotes, which converges to zero as K becomes large. Our bounds on the procedurally rational equilibrium provide managers with an easily-computed bound on the distance of said equilibrium from herding, aiding their operational planning. They also illustrate that if the service rates at the two stations are relatively far apart, then the outcome should be relatively close to herding. Second, for general route anecdotes, we show that as long as the probability of customers obtaining both types of anecdotes is high, our equilibrium will approximate herding. Hence, if customers are very likely to obtain anecdotes from both routes, then customers will herd. So, the firm can promote herding by facilitating information sharing among customers, perhaps by setting up a sharing platform to make it more likely that customers obtain anecdotes from both routes (like, e.g., the MyTSA app with crowdsourced waiting times for airport checkpoints: see Wang and Hu 2020).

2. Related Literature

As mentioned, Arlotto et al. (2019) find that in an open-routing service network, fully rational customers herd, and we also find that procedurally rational customers herd under some circumstances. Herding behavior has also been observed in the economics literature (see Smith and Sørensen 2020 for a recent example) as well as in other queueing-related settings (see Kremer and Debo 2016, Veeraraghavan and Debo 2011, among others). In Smith and Sørensen (2020), Bayesian customers use the actions taken by previous customers to update their beliefs about the utility of different actions. A string of customers choosing the same action influences the later customers to increase their quality belief for that action, which can lead to herding. In prior queueing studies including the two mentioned above, when customers choose between service providers and some have private signals about quality, the queue length conveys information about the quality of a service provider; this can lead to customers joining a longer queue to obtain higher quality service because the difference in quality can outweigh the increased waiting cost. Crucially, in all studies of herding behavior that we are aware of—apart from the open routing setting, that is—the driver of herding is *informational*. By contrast, in an open routing setting, herding is *strategic*. The more customers that choose a given route, the better that route becomes relative to other routes (see Arlotto et al. 2019); this is strikingly different from herding in other contexts, in which making the same decision as others either has no direct impact on utility (Smith and Sørensen 2020 and earlier studies of informational herding in economics) or actually *harms* the utility conditional on the quality level because it increases waiting time (Kremer and Debo 2016, Veeraraghavan and Debo 2011). Lastly, in the procedural rationality setting of this paper, customers reason based on anecdotes. While this might appear similar to reasoning based on the actions of others leading to informational herding, it is in fact fundamentally different. First, procedurally rational customers are not Bayesian but rather reason heuristically, and second, they decide based not only on the *actions* observed in their sample (as in the studies mentioned above) but also on the *consequences* of those actions, i.e., the realized system times.

More broadly, there is an extensive literature on strategic customer behavior in service systems, beginning with Naor (1969). Surveys can be found in Hassin and Haviv (2003) and Hassin (2016). Recent work on strategic customer behavior in service systems includes Yang and Debo (2019) and Cui et al. (2019), as well as several papers mentioned in Section 1 such as Wang and Hu (2020), which considers user-generated information sharing in a single-server queue with fully rational customers. There is also a burgeoning literature on modeling bounded rationality in operations management, which employs various customer behavioral models such as quantal choice and logit choice (Su 2008, Chen et al. 2012, Huang et al. 2013, Li et al. 2016), among others. A recent survey of this literature can be found in Ren and Huang (2018). In particular, an area of work that has been actively incorporating bounded rationality is the study of strategic queueing: Li et al. (2016) on quality-speed competition; Debo and Snitkovsky (2018) on tipping and social norms; Yang et al. (2018) on loss-averse customers; and Moon (2021) on customers choosing from a combinatorial set of paths in a network.

Several recent papers in operations study procedurally rational customers, e.g., Huang and Yu (2014) on opaque selling, Huang et al. (2017) on posterior price matching, and Ren et al. (2018) on join-balk decisions in a queueing system. Importantly, much of this work focuses on customers using anecdotal reasoning to infer *quality*. By contrast, in our model customers decide which route to follow through a service network, and they reason about their *waiting time* from choosing a given route. Another significant difference between this study and the existing economics and operations literature is that we have two types of anecdotes. In previous works on procedural rationality, it is conventional either to consider only a single type of anecdote or to assume that customers certainly have access to the same number of anecdotes (one or more) about every alternative. In either case, the composition of the sample (i.e., how many anecdotes from each alternative) is deterministic. This not only ignores the impact of the population proportion of alternatives chosen but also rules out the possibility that a customer may not receive any anecdote about one of the alternatives. By contrast, we propose a novel model which considers the availability of different types of anecdotes and how customers collect them, as well as how they choose in the event that one type of anecdote is missing entirely.

To our knowledge, the only previous work to assume that customers use anecdotal reasoning to infer waiting time in a service system is Huang and Chen (2015). They adopt the $S(1)$-reasoning framework, and customers make judgments about the expected waiting time from a single anecdote. By contrast, we demonstrate the nuanced role of procedural rationality in an open-routing service network, and as mentioned, in our context, customers receive anecdotes for more than one alternative.

3. The Model

We study a two-station service network with all customers present at the start of service. We label the stations as station S (Slow) and station F (Fast). Customers can choose which station to visit first, and every customer must visit both stations. A customer who visits station S (F) first is said to have chosen route SF (FS). The service rate at station S (μ_S) is assumed to be less than the service rate at station F (μ_F). After customers make their routing decisions, the SF customers—i.e., those who chose route SF—are sequenced uniformly at random to determine the queueing order at station S, and similarly for FS customers at station F. Upon completing service at the first station on her chosen route, a customer immediately joins the back of the other station's queue.

We consider a static routing game where players/customers choose route SF or FS. Our game setting with an open-routing service network resembles that of Arlotto et al. (2019). They consider a standard model with fully rational customers, where under the classical Nash equilibrium, each customer selects the route with the shortest expected system time given the strategy of all others. By contrast, we propose a model where customers are not aware of the strategies of other customers, nor any system parameters for that matter. Instead, customers are procedurally rational and select routes based on the anecdotes they sample from the population. Finally, we consider a procedurally rational equilibrium in a game with a large number of players, where each player's decision is consistent with the population's procedurally rational behavior.

Next, we offer an alternative interpretation of the procedurally rational equilibrium in our static open routing game model, in the form of a repeated game. Consider a repeated game where on each day, a large number of customers, approximated as fluid in our analysis, participate in the open-routing game. When customers make decisions on a given day, they rely on the anecdotes they have collected, which are sampled from the experiences of the customer population on a previous day. In this repeated game, a steady state is an outcome in which the customer population has reached a stable and consistent pattern of behavior, which emerges through interactions between the current customers and the past customers via anecdotes. Note that a procedurally rational equilibrium in the static game is equivalent to a steady state in the repeated game.

To further validate the relationship between our static game's procedurally rational equilibrium and the repeated interaction interpretation, we conduct numerical experiments (details in Section 6). In our experiments, customers are discrete (hence non-fluid), and on each day, they act based on anecdotes collected from past customers. Our observations demonstrate that the simulation converges to an outcome that closely resembles the equilibrium of our static fluid model. Our numerical findings thus highlight the applicability of the intuitions and understandings derived from studying the static routing model to a scenario where customers behave anecdotally based on past customers' experiences.

Next, we formally describe the decision process of procedurally rational customers. Similar to the procedural rationality literature, we assume each customer is choosing her route based on her estimate of the system time for each route, which depends on other customers' actions. A given customer starts with a prior estimate S_0 for the expected system time for the SF route (and similarly F_0 for the FS route). We model S_0 and F_0 as normal random variables, with means equal to the true expected system times corresponding to the SF and FS routes and standard deviation of σ. In our model, S_0 (or F_0) is used as the estimated expected system time only if the customer does not obtain any anecdotes for route SF (or FS). However, if the customer does receive anecdotes $x_1, \ldots, x_k$ for a route with $k \geq 1$, the estimated expected system time is updated to the average, $(x_1 + \ldots + x_k)/k$, which does not include the prior. After estimating the system times using her anecdotes, the customer chooses a route based on her estimates. Note that when σ, the noise of the prior estimate, is large, a customer with a sample containing only one type of anecdote will select either route SF or FS with the probability of approximately 1/2 because with very large σ, the variation in the prior for the missing route far exceeds the variation in the anecdotes for the other route.

One could propose an expanded model where the estimated expected system time is calculated as a convex combination of $(x_1 + \ldots + x_k)/k$ and the prior, similar to Huang et al. (2017). However, we adopt the more parsimonious model, as a simplifying assumption for the analysis and because the effects on equilibrium outcomes from the expanded model are fairly predictable. Our parsimonious model is equivalent to the weight for the prior being zero when anecdotes for the route are obtained. If a higher weight is assigned to the prior, customer behavior would be propelled towards or away from fully rational behavior depending on the value of σ. If σ is small, a higher weight on the prior nudges customer behavior towards full rationality, and conversely if σ is large (see Appendix M for details).

We will explore two model variations pertaining to how customers sample anecdotes. The first variation, *random route anecdotes*, assumes that each customer samples a fixed total number K of anecdotes randomly from the entire population.[3] The second variation, *general route anecdotes*, assumes that K_S, the number of anecdotes from the SF route, and K_F, the number of anecdotes from the FS route, follow an arbitrary joint probability distribution. The general route anecdotes variation allows us to model more nuanced customer sampling behaviors. For instance, customers, in order to make a more informed decision, may be biased toward drawing/receiving anecdotes from the sub-population of the missing type. Appendix E gives a concrete example of such a sampling process.

Given a customer sampling process, we study a model where customers are depicted as fluid. In our fluid model, a certain volume (normalized to 1) of fluid must be processed at both stations. Station S processes fluid at rate μ_S, and station F processes fluid at rate $\mu_F \geq \mu_S$. Appendix B gives technical details on the system evolution and customer system times in the fluid model.

3.1. Response Function under Fluid Approximation

Next, we derive the fluid response function for our procedurally rational model. Let $K_S(\alpha)$ ($K_F(\alpha)$) represent the (random) number of anecdotes a customer receives from route SF (FS), given a fraction α of (fluid) customers on route SF. We denote the collections of SF and FS anecdotes by $\boldsymbol{S} := (S_1, \ldots, S_{K_S(\alpha)})$ and $\boldsymbol{F} := (F_1, \ldots, F_{K_F(\alpha)})$, where one of these sets may be empty. In our model, we assume that each anecdote of route SF (or FS) represents the system time of a customer randomly and independently drawn from all the customers in route SF (or FS).

Define $\hat{S}(\alpha)$ ($\hat{F}(\alpha)$) as a customer's estimate of the system time on route SF (FS) after obtaining her sample, given that the sample is drawn from a population with an SF fraction of α. Recall that the customer uses her prior as her estimate whenever no anecdote for the route is obtained. Thus, the estimates $\hat{S}(\alpha)$ and $\hat{F}(\alpha)$ are given by

$$\hat{S}(\alpha) = \begin{cases} \sum_{i=1}^{K_S(\alpha)} S_i / K_S(\alpha) & \text{if } K_S(\alpha) \geq 1, \\ S_0 & \text{if } K_S(\alpha) = 0. \end{cases} \quad \text{and} \quad \hat{F}(\alpha) = \begin{cases} \sum_{j=1}^{K_F(\alpha)} F_j / K_F(\alpha) & \text{if } K_F(\alpha) \geq 1, \\ F_0 & \text{if } K_F(\alpha) = 0. \end{cases}$$

Note that $\hat{S}(\alpha)$ and $\hat{F}(\alpha)$ have three sources of randomness: the random prior point estimate (S_0 or F_0), the number of anecdotes ($K_S(\alpha)$ or $K_F(\alpha)$), and the random anecdotes themselves (S_i or F_i). Moreover, the distributions of the prior point estimate and of the anecdotes depend on α.

In the fluid regime, the fraction of customers who choose route SF is the same as $\mathsf{Pr}[\hat{S}(\alpha) \leq \hat{F}(\alpha)]$, and it is this probability that we call the *response function*, which we denote by π. We seek an equilibrium point of this response function, namely an SF fraction α such that

$$\pi(\alpha) := \mathsf{Pr}[\hat{S}(\alpha) \leq \hat{F}(\alpha)] = \alpha. \tag{1}$$

[3] As we consider a large population (fluid) model, anecdotes sampled from the population are assumed to be independent and identically distributed (i.i.d.).

Hereafter, we first analyze the response function of the random route anecdotes model, where each anecdote is randomly sampled from the entire population, and then we study the general route anecdotes model to obtain additional insights. Moreover, although our model assumes that S_0 and F_0 for each customer are drawn from a normal distribution centered at the true expected system time, our results remain unchanged as long as the prior is centered anywhere within the *support* of the system time (see Appendix L for details).

4. Equilibrium Analysis

Let $\pi^{FR}(\alpha)$ denote the (fully rational) best response function, i.e., the fraction of customers that will choose route SF by playing their best response on their expected system time for both routes given that a fraction α of customers are choosing SF.[4] As a result, α^{FR} is a Nash equilibrium of the fully rational model if and only if $\pi^{FR}(\alpha^{FR}) = \alpha^{FR}$. Arlotto et al. (2019) proved that all customers herding on route SF (or FS when $2\mu_S > \mu_F$) is a Nash equilibrium in the routing game with discrete customers. The analogous result also holds with a continuum of customers, as the following claim demonstrates. The proof of the claim appears in Appendix C.

Claim 1. *For any μ_S and μ_F, we have $\pi^{FR}(1) = 1$. If $2\mu_S \geq \mu_F$, we also have that $\pi^{FR}(0) = 0$.*

Remark 1. In general, when $2\mu_S \geq \mu_F$, the fully rational model may have a third Nash equilibrium in addition to the two herding equilibria. Throughout, we will focus on the herding equilibrium at SF ($\alpha^{FR} = 1$) for two reasons. First, herding on SF is a Nash equilibrium for all $\mu_S \leq \mu_F$, whereas the other Nash equilibria only hold when $\mu_S \leq \mu_F \leq 2\mu_S$. Second, among all Nash equilibria, the herding equilibrium at SF has the smallest cumulative system time over all customers (see Section 5). Thus, it is in everyone's best interest to coordinate the herding equilibrium at SF.

To ensure meaningful comparisons with the Nash equilibrium with $\alpha^{FR} = 1$, we focus on the largest equilibrium in our procedurally rational model, i.e., that closest to 1. Although our model can have multiple equilibria, we treat the largest procedurally rational equilibrium as focal because the service provider may be able to facilitate among the equilibria, and it is in their interest to facilitate the closest equilibrium to $\alpha^{FR} = 1$, which minimizes the cumulative system time for all customers.

Definition 1. Given a procedurally rational model, we define $\alpha^* := \max\{\alpha : \pi(\alpha) = \alpha\}$, where π is the response function given in (1).

4.1. Procedurally Rational Equilibrium with Random Route Anecdotes

In this subsection, we focus on the procedurally rational model with K *random route anecdotes* and compare it with the equilibrium under the fully rational model. Recall that under the random route anecdotes model, each anecdote is randomly sampled from the entire population. Thus, $K_S(\alpha)$ follows

[4] If SF and FS have the same expected system time for SF fraction $0 < \alpha < 1$, either route (SF or FS) is a best response. In that case, we can assume that customers choose SF with probability α, making α a Nash equilibrium.

a binomial distribution with K trials and success probability α, while $K_F(\alpha) = K - K_S(\alpha)$. Therefore, the response function, π, can be expressed as

$$\begin{aligned}\pi(\alpha) = (1-\alpha)^K \mathsf{Pr}\left[S_0 \le \frac{\sum_{j=1}^K F_j}{K}\right] + \alpha^K \mathsf{Pr}\left[\frac{\sum_{j=1}^K S_j}{K} \le F_0\right] \\ + \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S} (1-\alpha)^{K-k_S} \cdot \mathsf{Pr}\left[\frac{\sum_{i=1}^{k_S} S_i}{k_S} \le \frac{\sum_{j=1}^{K-k_S} F_j}{K-k_S}\right].\end{aligned} \tag{2}$$

Recall that the distributions of S_j and F_j for $0 \le j \le K$ depend on the SF fraction α. To highlight this dependence, we sometimes use $S_j(\alpha)$ and $F_j(\alpha)$ in place of S_j and F_j. In addition, we define

$$\begin{aligned}\pi_0(\alpha) = \mathsf{Pr}\left[S_0(\alpha) \le \frac{\sum_{j=1}^K F_j(\alpha)}{K}\right], \quad \pi_K(\alpha) = \mathsf{Pr}\left[\frac{\sum_{j=1}^K S_j(\alpha)}{K} \le F_0(\alpha)\right], \\ \text{and } \pi_k(\alpha) = \mathsf{Pr}\left[\frac{\sum_{i=1}^k S_i(\alpha)}{k} \le \frac{\sum_{j=1}^{K-k} F_j(\alpha)}{K-k}\right], \forall k = 1, \ldots K-1.\end{aligned} \tag{3}$$

With the definition of π_k, we can simplify (2) as

$$\pi(\alpha) = (1-\alpha)^K \pi_0(\alpha) + \alpha^K \pi_K(\alpha) + \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S} (1-\alpha)^{K-k_S} \pi_{k_S}(\alpha). \tag{4}$$

We will compare the outcomes of the fully and procedurally rational models via the difference between the herding Nash equilibrium ($\alpha^{FR} = 1$) and α^*. Intuitively, for each $0 < \alpha < 1$, as the number of anecdotes collected by each customer (denoted by K) becomes large, each customer should get nearly perfect inference on the true expected system time from choosing route SF or FS, given that α fraction of customers are choosing SF. Therefore, the outcome of our procedurally rational model should approach that of the rational model for large K. Our next proposition verifies this intuition, by showing that the largest equilibrium outcome of our model converges to the herding (Nash) equilibrium as K approaches ∞.

PROPOSITION 1 **(Equilibrium For Large Sample Size)**. *Consider a sequence of procedurally rational models with changing K and fixed σ. Then we have* $\lim_{K\to\infty} \alpha^* = 1$.

In addition to having the same equilibrium outcomes, we note that the response function under our procedurally rational model converges to that of the fully rational model in the interval $(1/2, 1)$ as K approaches infinity. This is stated as Corollary 1, which follows from the proof of Proposition 1.

COROLLARY 1. *For any fixed $\alpha \in (1/2, 1)$, we have* $\lim_{K\to\infty} \pi(\alpha) = \pi^{FR}(\alpha) = 1$, *where π is the response function given in* (2).

Thus, as customers receive more anecdotes, their estimates of the system time approach the true mean; while not altogether surprising, this finding is valuable in that it establishes the connection between procedural rationality and full rationality. Moreover, as we will demonstrate in Section 5, the cumulative system time for all customers is minimized at $\alpha = 1$. Consequently, it is socially optimal for customers to herd. Hence, our result demonstrates that if all customers were able to obtain more anecdotes, i.e., increasing K, it would enhance the overall customer experience.

We have seen that customer behavior under procedural rationality converges to the socially optimal herding equilibrium as the sample size grows to infinity, but in practice, access to a large number of anecdotes can come at a cost to people, as they may be simply overwhelmed by information or have limited memory (see, e.g., Tong and Feiler 2017 for more on cognitive limitations related to sampling). Also, the number of each type of anecdote they receive may not be binomial if the sampling is biased for some reason (e.g., to obtain a missing type of anecdote). Therefore, we next seek to understand the largest equilibria for procedurally rational customers given that the sample size, K, is finite.

At first glance, equilibria associated with finite K appear to be substantially more intricate due to the complex convolution of random variables within $\hat{S}(\alpha)$ and $\hat{F}(\alpha)$. To overcome the complication of analyzing convolutions in the context of finite K, we make a critical observation: around the equilibrium, the support intervals for the random system time draws on routes SF and FS intersect only at a single point. As a result, any system time drawn from route SF will always be shorter than that from route FS with probability one.

Property 1. *Consider a customer drawing a sample of anecdotes from a system with $\alpha \in (\mu_S/\mu_F, 1)$. Let $S(\alpha)$ and $F(\alpha)$ denote the distributions of an anecdote from the SF and FS routes, respectively. Similarly, let* $\mathrm{supp}(S(\alpha))$ *and* $\mathrm{supp}(F(\alpha))$ *represent their respective supports. We have*

$$\sup\{\mathrm{supp}(S(\alpha)\} = \inf\{\mathrm{supp}(F(\alpha))\} = \alpha/\mu_S, \tag{5}$$

$$\text{and } \pi_k(\alpha) = \mathsf{Pr}\Big[\frac{\sum_{i=1}^k S_i(\alpha)}{k} \leq \frac{\sum_{j=1}^{K-k} F_j(\alpha)}{K-k}\Big] = 1 \quad \text{for any } 1 \leq k \leq K-1. \tag{6}$$

The property above (proof in Appendix B) shows that when $\alpha \in (\mu_S/\mu_F, 1)$, there is a limit on the amount of noise in the anecdotes a customer receives. More importantly, the noise is bounded in such a way that whenever a customer draws anecdotes from both routes, she will correctly conclude that route SF has the smaller expected system time with probability one. This property plays a pivotal role in our subsequent analysis of the procedurally rational equilibrium with a fixed and finite K.

Observe that for a fixed K, the procedurally rational equilibrium is influenced not only by K, but also by σ, the noise associated with the priors. Intuitively, as σ increases, α^* deviates further from herding as the noise introduced by the priors leads to a greater fraction of customers incorrectly identifying the route with the shorter system time. Nevertheless, we find by applying Property 1 that the impact of σ diminishes as K increases as far as the α^* is concerned. Following this, we present a result with a straightforward-to-calculate lower limit for α^* that applies for all values of $\sigma \geq 0$.

Proposition 2 **(Equilibrium Bound and Convergence Rate)**. *For any $\sigma \geq 0$:*

- *if $\mu_S/\mu_F < 1/\sqrt{2} \approx 0.707$, then $\alpha^* \geq 1/\sqrt{2}$ for any positive integer K,*
- *for any integers $K \geq (\frac{\mu_F}{\mu_F - \mu_S})^2$, we have $\alpha^* \geq 1 - 1/\sqrt{K}$.*

Proposition 2 provides a closed-form lower bound for α^*, which, importantly, holds for non-limiting cases (i.e., it does not require extreme values of K or σ). While Proposition 2 may seem to extend Proposition 1, it requires a very different approach as it is difficult to characterize the distributions of convolutions of the SF and FS anecdotes. Instead, our proof is critically reliant on the application of Property 1. In addition to providing a straightforward and closed-form lower limit for α^*, Proposition 2 also establishes a convergence rate of $O(1/\sqrt{K})$ for α^* as K increases.

We next consider the effect of σ on the procedurally rational equilibrium. Note that when a sample contains only anecdotes from route SF (or alternatively, FS), the customer may incorrectly infer that route SF has the longer expected system time due to the noise in the distribution of $F_0(\alpha)$ (or $S_0(\alpha)$) measured by σ. Intuitively, the probability that a customer identifies the wrong route with the shorter system time decreases as σ becomes small, which drives α^* closer to herding. The convergence to herding as σ decreases to zero is formalized in the next proposition.

Proposition 3 **(Small Prior Noise Is Sufficient for Herding)**. *For any fixed K, as σ converges to 0, we have that α^* converges to 1.*

Proposition 3 establishes the following interesting insight. Suppose the prior estimate has negligible noise ($\sigma = 0$); then the equilibrium outcome will be close to herding, *even if the customers ignore the prior in the event that they receive anecdotes for a route.*

To provide intuition, Figure 1 plots the response function (computed by Monte Carlo simulation) for various values of σ and K. The 45-degree line $f(\alpha) = \alpha$ is plotted for reference, and equilibrium occurs where the response function intersects this line. As implied by Propositions 2 and 3, the figure illustrates that α^* converges to herding (i.e., to $\alpha = 1$) as either (i) the noise of the prior becomes small (small σ), or (ii) the sample size K becomes large. We also observe that α^* is close to herding across a wide range of K and σ values. For instance, with $K = 2$ and $\sigma = 0.25$, α^* is already at 0.96.

4.2. General Route Anecdotes

Thus far, our exploration has focused on a model with procedurally rational customers, each of whom draws exactly K random route anecdotes. While this model does encapsulate some key aspects of boundedly rational customer behavior, it remains a simplification of the complex reality. For example, it misses the feature that customers in the real world may draw a random number of anecdotes, and the distribution for the number of each type of anecdote they receive may deviate from a binomial distribution. Consequently, we aim to understand the behavior of customers with *general route*

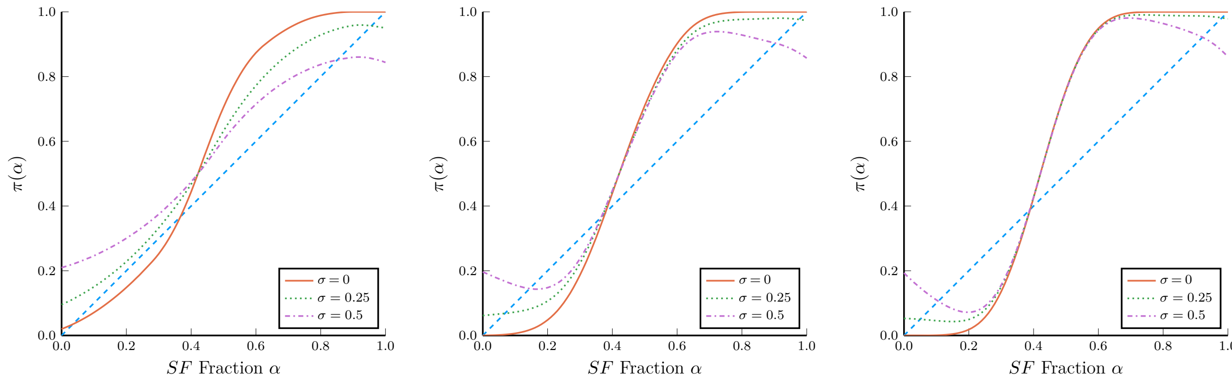

(a) $K = 2$ (b) $K = 5$ (c) $K = 8$

Figure 1 Procedurally rational response function $\pi(\alpha)$ for different σ and K ($\mu_S = .9$, $\mu_F = 1$).

anecdotes, where K_S, the number of anecdotes from the SF route, and K_F, the number of anecdotes from the FS route, may follow an arbitrary joint probability distribution.

Similar to Proposition 1, intuitively, as K_S and K_F go to infinity, customers' behavior tends to align closely with full rationality and α^* converges to herding. When K_S and K_F are small, an uninformed observer may expect that sampling error would cause customers who reason via anecdotes to deviate in their behavior from the fully rational norm of herding, with this deviation being more prominent when K_S and K_F are small. However, our next proposition shows that the largest equilibrium converges to herding (i.e., to $\alpha = 1$), as long as the probability of a customer obtaining both types of anecdotes approaches 1. Like Proposition 2 and Proposition 3, the proof of the next proposition hinges on Property 1, which remains applicable in the general route anecdotes case because the sampling process does not affect the support of the system times for any given α.[5]

PROPOSITION 4 **(Obtaining Both Types of Anecdotes Is Sufficient for Herding)**. *Consider a procedurally rational model with general route anecdotes parameterized by a quantity β. Let $\gamma(\alpha, \beta)$ be the probability of a customer obtaining at least one anecdote from each route (SF and FS). Suppose that (i) the response function $\pi^\beta(\alpha)$ is continuous in α for any β; and (ii) $\lim_{\beta\to\infty} \gamma(\alpha, \beta) = 1$ for any $\alpha \in (0, 1)$. Then, we have $\lim_{\beta\to\infty} \alpha^* = 1$, where α^* is the largest equilibrium when the sampling parameter is β.*

In other words, whether the sample size is small or large, as long as the likelihood of obtaining anecdotes for both routes is high, the procedurally rational equilibrium comes to resemble the fully rational equilibrium. This finding implies the following important and encouraging insight for managers: if customers are likely to obtain information about both routes, then increasing the

[5] Formally, Equation (6) of Property 1 continues to hold for any pair of k_S and k_F with $k_S, k_F \geq 1$.

number of anecdotes—which can be thought of as customers becoming more sophisticated in their reasoning—is not necessary to achieve a good customer experience.

We have intentionally defined the general route anecdotes sampling process abstractly to retain the most generality. However, for additional concreteness, in Appendix E we provide a specific example of a general route anecdotes sampling process. In this process, customers' anecdotes are biased toward the missing type of anecdote until they obtain both types, and the degree of bias is related to a *discernibility* parameter. This parameter is a simplified abstraction of the factors other than the population proportion of routes that influence the probability of obtaining the missing type of anecdote, e.g., the size and activeness of the online community in the social media example in the introduction. This sampling process with discernibility satisfies the conditions of Proposition 4, so as discernibility increases, α^* converges to herding. For details, refer to Appendix E.

5. Cumulative System Time

To measure how well the system performs—and by extension, how satisfied customers are likely to be—we use the *cumulative system time*, i.e., the integral of the total system time experienced by customers in each position in the queues (see Proposition D.1 in Appendix D for a closed-form expression). Arlotto et al. (2019) report that herding achieves excellent performance with respect to cumulative system time in the discrete setting. The next proposition verifies that in the fluid case also, herding performs extremely well and is, in fact, optimal.

Proposition 5 **(Herding Is Socially Optimal)**. *The cumulative system time is minimized when customers herd, i.e., for fixed μ_S and μ_F, we have* $\arg\min_{\alpha\in[0,1]} D(\alpha,\mu_S,\mu_F)=\{0,1\}$.

We also compare the cumulative system time at the procedurally rational equilibrium against that at herding for different values of the service rate ratio μ_S/μ_F and the prior noise σ. When the service rates are far apart, the probability of an SF draw being worse than an FS draw is small for almost any α, so the equilibrium under procedural rationality is close to herding and the cumulative system time is thus near optimal. When the service rates are close together, the noise in the prior plays a larger role. For larger σ, there can be meaningful performance loss compared to herding, but this diminishes as σ decreases (see Appendix D.2 for details).

6. Accuracy of the Fluid Model

With discrete customers, we now numerically study the evolution of procedurally rational routing decisions in a repeated setting. Specifically, we suppose that new customers enter to play the game in every period, and their anecdotes are drawn from customers in the previous period. We again let $K=2$. We initialize $\alpha=.75$ in the first round, midway between 1/2 and 1. For a range of service rates μ_S (we fix $\mu_F=1$ in our simulations) and prior noise parameters σ (which we scale by N for

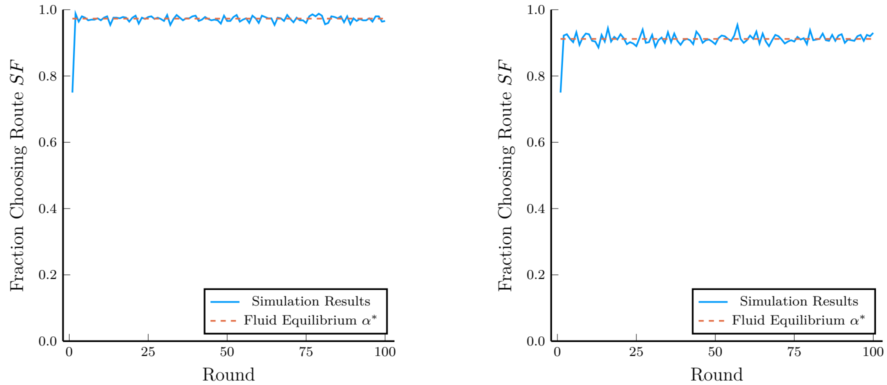

(a) $\sigma = .25$ (b) $\sigma = .5$

Figure 2 Simulated evolution of discrete system ($N = 500, \mu_S = .7, \mu_F = 1$).

appropriate comparison with the mass 1 fluid model), we simulated systems with $N = 500$ customers for 100 rounds each, replicating each system 10 times.

First, we observe that, even for systems with only 500 customers, our analytical results for the fluid system are very good predictors of the routing decisions of procedurally rational customers. Table A.1 in Appendix A reports sample statistics for the fraction of SF customers in rounds 75-100, across replications for each parameter combination. We observe that play tends to hover very closely around the fluid equilibrium. The medians are very close to α^*, and with the exception of the extreme case with equal service rates, the first and third quartiles are between 0 and .04 above or below α^*. In addition, for given service rates, we observe that the interquartile range tends to be larger for larger values of σ, which is to be expected as higher σ entails more variation in the prior.

Figure 2 shows typical sample paths of the discrete system. We observe a similar effect in these plots, namely that play quickly approaches α^* and then fluctuates in a relatively narrow band around it. Comparing the left and right panels of the figure, we can see that smaller σ brings the largest equilibrium closer to herding, consistent with Proposition 2. In addition, for the right panel with larger σ, the SF fraction fluctuates more, consistent with our observation that the interquartile range increases with σ. Overall, our results with discrete customers reflect similar behavior to that in the fluid model. In addition, in Appendix F, we study systems with heterogeneity in the sample size and the prior noise, and we find that customer behavior closely resembles that under our main model.

7. Conclusion

We study an open-routing service network with two stations and self-interested, procedurally rational customers who make decisions about which route to take through the network based on anecdotes. In addition to the managerial insights described next, we make a theoretical contribution to the procedural rationality literature by explicitly modeling the sampling process by which customers obtain anecdotes, as well as what customers do if they are missing one type of anecdote.

In the fully rational counterpart to our model, customers herd, i.e., they all take the same route through the network; this outcome is also socially optimal in terms of cumulative system time. By comparing outcomes with procedurally rational customers to the fully rational outcome, we uncover three managerial insights. First, we obtain closed-form bounds on the distance between the procedurally rational and fully rational equilibrium, aiding operational planning and showing the rate of convergence to the fully rational outcome as the sample size grows. Second, if customers obtain anecdotes of both actions with high probability, then the equilibrium will approximate the fully rational outcome, *despite the sampling error inherent to procedural rationality.* Finally, if the noise in the prior point estimate is small, then the procedurally rational equilibrium approximates the fully rational equilibrium, *even if customers ignore the prior when they have anecdotes.*

These insights reveal two different ways that managers with procedurally rational customers can facilitate herding. First, managers can promote herding by facilitating information sharing among customers so that they are more likely to obtain anecdotes about both routes. Surprisingly, even if each customer's sample is quite small, limited but diverse information encompassing both alternatives can still significantly impact customer decisions towards the fully rational and socially optimal outcome of herding. Second, if managers can influence the prior point estimates by providing information with minimal noise, then the equilibrium will again approximate herding. Importantly, this information need not even be very accurate; even if the prior is not centered at the true expected system time, if the noise (i.e., standard deviation) in the prior is small, then as long as the likelihood of the prior estimate being in the *support* of the system time is high, the equilibrium will approximate herding.

We remark, however, that the fully rational outcome may not always be the goal. Many examples in the literature reveal sub-optimal outcomes due to selfish customer behavior, and the cost of such behavior has even been formalized in the price of anarchy (see, e.g., Roughgarden and Tardos 2002). So, in settings where the fully rational outcome is undesirable, the preferred managerial interventions might be strikingly different from what we have identified in this work. But regardless of whether the fully rational outcome is desirable in a particular context, we believe that our approach that directly models the sampling process improves the fidelity of the procedural rationality framework, and it may be especially useful to researchers studying procedurally rational customers in service settings.

Acknowledgments

The authors thank department editor Michael Pinedo, as well as the anonymous senior editor and referees, for many helpful suggestions that led to a significantly improved paper.

References

Akan, M., Ata, B., and Olsen, T. (2012). Congestion-based lead-time quotation for heterogenous customers with convex-concave delay costs: Optimality of a cost-balancing policy based on convex hull functions. *Operations Research*, 60(6):1505–1519.

Arlotto, A., Frazelle, A. E., and Wei, Y. (2019). Strategic open routing in service networks. *Management Science*, 65(2):735–750.

Baron, O., Berman, O., Krass, D., and Wang, J. (2016). Strategic idleness and dynamic scheduling in an open-shop service network: Case study and analysis. *Manufacturing & Service Operations Management*, 19(1):52–71.

Chen, Y., Su, X., and Zhao, X. (2012). Modeling bounded rationality in capacity allocation games with the quantal response equilibrium. *Management Science*, 58(10):1952–1962.

Cui, S., Su, X., and Veeraraghavan, S. (2019). A model of rational retrials in queues. *Operations Research*, 67(6):1699–1718.

Debo, L. and Snitkovsky, R. I. (2018). Tipping in service systems: The role of a social norm. *Working Paper, Dartmouth College, Hanover, NH.*

Hassin, R. (2016). *Rational Queueing*. Chapman and Hall/CRC.

Hassin, R. and Haviv, M. (2003). *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*. Springer Science & Business Media.

Honnappa, H. and Jain, R. (2015). Strategic arrivals into queueing networks: The network concert queueing game. *Operations Research*, 63(1):247–259.

Huang, T., Allon, G., and Bassamboo, A. (2013). Bounded rationality in service systems. *Manufacturing & Service Operations Management*, 15(2):263–279.

Huang, T. and Chen, Y.-J. (2015). Service systems with experience-based anecdotal reasoning customers. *Production and Operations Management*, 24(5):778–790.

Huang, T., Yin, Z., and Chen, Y.-J. (2017). Managing posterior price matching: The role of customer boundedly rational expectations. *Manufacturing & Service Operations Management*, 19(3):385–402.

Huang, T. and Yu, Y. (2014). Sell probabilistic goods? A behavioral explanation for opaque selling. *Marketing Science*, 33(5):743–759.

Juslin, P., Winman, A., and Hansson, P. (2007). The naïve intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review*, 114(3):678.

Kremer, M. and Debo, L. (2016). Inferring quality from wait time. *Management Science*, 62(10):3023–3038.

Li, X., Guo, P., and Lian, Z. (2016). Quality-speed competition in customer-intensive services with boundedly rational customers. *Production and Operations Management*, 25(11):1885–1901.

Moon, K. (2021). When customers choose routes in service networks: Estimating complementarities and congestion effects using machine learning. *Working Paper, University of Pennsylvania, Philadelphia, PA. Available at SSRN 3819117.*

Naor, P. (1969). The regulation of queue size by levying tolls. *Econometrica: Journal of the Econometric Society*, 37(1):15–24.

Osborne, M. J. and Rubinstein, A. (1998). Games with procedurally rational players. *American Economic Review*, 88(4):834–847.

Parlaktürk, A. K. and Kumar, S. (2004). Self-interested routing in queueing networks. *Management Science*, 50(7):949–966.

Ren, H. and Huang, T. (2018). Modeling customer bounded rationality in operations management: A review and research opportunities. *Computers & Operations Research*, 91:48–58.

Ren, H., Huang, T., and Arifoglu, K. (2018). Managing service systems with unknown quality and customer anecdotal reasoning. *Production and Operations Management*, 27(6):1038–1051.

Roughgarden, T. and Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM*, 49(2):236–259.

Shtrichman, O., Ben-Haim, R., and Pollatschek, M. A. (2001). Using simulation to increase efficiency in an army recruitment office. *Interfaces*, 31(4):61–70.

Smith, L. and Sørensen, P. N. (2020). Rational social learning with random sampling. *Working Paper, University of Wisconsin-Madison, Madison, WI.*

Spiegler, R. (2006). The market for quacks. *The Review of Economic Studies*, 73(4):1113–1131.

Su, X. (2008). Bounded rationality in newsvendor models. *Manufacturing & Service Operations Management*, 10(4):566–589.

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

Tversky, A. and Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76(2):105–110.

Veeraraghavan, S. K. and Debo, L. G. (2011). Herding in queues with waiting costs: Rationality and regret. *Manufacturing & Service Operations Management*, 13(3):329–346.

Wang, J. and Hu, M. (2020). Efficient inaccuracy: User-generated information sharing in a queue. *Management Science*, 66(10):4648–4666.

Wu, S., Liu, Q., and Zhang, R. Q. (2015). The reference effects on a retailer's dynamic pricing and inventory strategies with strategic consumers. *Operations Research*, 63(6):1320–1335.

Yang, L. and Debo, L. (2019). Referral priority program: Leveraging social ties via operational incentives. *Management Science*, 65(5):2231–2248.

Yang, L., Guo, P., and Wang, Y. (2018). Service pricing with loss-averse customers. *Operations Research*, 66(3):761–777.

Online Appendix to:

Service Networks with Open Routing and Procedurally Rational Customers

Andrew E. Frazelle, Tingliang Huang, and Yehua Wei

This appendix is divided into sections. We first provide an additional table in Appendix A. Then, we give technical details of our fluid model in Appendix B, including the proof of Property 1. In Appendix C, we prove Claim 1. Appendix D provides details and a numerical study of the cumulative system time, including the closed-form expression and its proof (Proposition D.1). We discuss a concrete example of the general route anecdotes sampling process in Appendix E. In Appendix F, we consider systems with heterogeneous customers. In Appendices G-K, we provide the proofs of the propositions in the main body, including any auxiliary results or corollaries, with a section for each proposition. Finally, in Appendices L and M, we discuss respectively the cases where the prior point estimate is not centered at the true expected system time, and where customers assign positive weight to the prior for a route even when they have anecdotes from that route.

A. Additional Table

Table A.1 SF **fraction in rounds 75-100 for Section 6 simulations (** $\mu_F = 1$, $N = 500$, **100 rounds, 10 replications)**

	μ_S					
σ	0.5	0.6	0.7	0.8	0.9	1
0.25	0.986	0.980	0.972	0.966	0.956	0.914
0.5	0.950	0.930	0.912	0.892	0.850	0.727
0.75	0.904	0.880	0.856	0.830	0.784	0.638
1	0.866	0.842	0.820	0.798	0.748	0.546

(a) Median

	μ_S					
σ	0.5	0.6	0.7	0.8	0.9	1
0.25	0.986	0.980	0.973	0.965	0.957	0.914
0.5	0.948	0.931	0.912	0.891	0.852	0.723
0.75	0.904	0.879	0.856	0.831	0.780	0.643
1	0.866	0.841	0.819	0.795	0.746	0.602

(b) Equilibrium α^*

	μ_S					
σ	0.5	0.6	0.7	0.8	0.9	1
0.25	0.982	0.976	0.968	0.960	0.948	0.904
0.5	0.940	0.922	0.904	0.882	0.838	0.712
0.75	0.894	0.868	0.846	0.819	0.768	0.609
1	0.855	0.828	0.806	0.784	0.736	0.450

(c) First Quartile

	μ_S					
σ	0.5	0.6	0.7	0.8	0.9	1
0.25	0.990	0.984	0.976	0.970	0.960	0.924
0.5	0.954	0.938	0.920	0.900	0.862	0.748
0.75	0.914	0.890	0.866	0.844	0.795	0.664
1	0.874	0.852	0.832	0.808	0.760	0.602

(d) Third Quartile

B. Fluid Model Details

Given its service rate μ_S, in an elapsed time of length ℓ, station S is capable of processing a volume $\mu_S \ell$ of fluid. Correspondingly, to process a volume v of fluid at station S requires a length of time equal to

$$T_S(v) = \frac{v}{\mu_S}. \tag{B.1}$$

We similarly have that in an elapsed time of length ℓ, station F can process a volume $\mu_F \ell$ of fluid. To processes a volume v of fluid at station S requires a length of time equal to

$$T_F(v) = \frac{v}{\mu_F}. \tag{B.2}$$

These expressions can be thought of as approximating the limiting case for a discrete system as the number of customers N grows large, if we let the service rates grow proportionally with the number of customers in the system.

Recall that we use α to denote the fraction α of customers that chooses route SF. Denote by $Q_S(\alpha;\ell)$ the amount of fluid waiting in the queue (or "buffer") at station S given an SF fraction of α, an elapsed time ℓ after the system begins operating. Define $Q_F(\alpha;\ell)$ similarly for station F. Thus, we have $Q_S(\alpha;0)=\alpha$, and $Q_F(\alpha;0)=1-\alpha$. Letting $[x]^+=\max\{x,0\}$, the amount of fluid $Q_F(\alpha;\ell)$ in station F's buffer after the system has been operating for an elapsed time ℓ is equal to

$$\begin{aligned} Q_F(\alpha;\ell) &:= \Big[Q_F(\alpha;0)+\min\{Q_S(\alpha;0),\mu_S\ell\}-\mu_F\ell\Big]^+ \\ &= \Big[1-\alpha+\min\{\alpha,\mu_S\ell\}-\mu_F\ell\Big]^+. \end{aligned} \tag{B.3}$$

This relation takes the positive part of a simple balance equation: after a time ℓ, the amount in the buffer at station F is equal to the initial quantity, plus the fluid that has arrived from station S, minus the fluid that has been processed. Because fluid arrives to station F at a slower rate than it is processed, the quantity in the buffer will be strictly decreasing in ℓ until it hits zero, where it remains. We can similarly express the amount of fluid $Q_S(\alpha;\ell)$ in the station S buffer after a time ℓ has elapsed by

$$\begin{aligned} Q_S(\alpha;\ell) &:= \Big[Q_S(\alpha;0)+\min\{Q_F(\alpha;0),\mu_F\ell\}-\mu_S\ell\Big]^+ \\ &= \Big[\alpha+\min\{1-\alpha,\mu_F\ell\}-\mu_S\ell\Big]^+. \end{aligned} \tag{B.4}$$

Note that the buffer at station S will first increase with time because fluid arrives to station S faster than it is processed. This increase will continue until the entire volume $1-\alpha$ of FS customers has departed station F to join the queue at station S, after which the station S buffer will shrink at rate μ_S until it empties.

Let y_S be a possible starting position in the buffer at station S, where $0\le y_S\le\alpha$, and let y_F be a possible starting position in the queue at station F, where $0\le y_F\le 1-\alpha$. We denote by $\mathcal{S}(\alpha;y_S)$ the total system time for the infinitesimal customer starting in position y_S in the station S queue, given a total SF fraction α.

Lemma B.1 **(System Time for *SF* Customers)**. *The function $\mathcal{S}(\alpha;y_S)$ can be expressed as*

$$\mathcal{S}(\alpha;y_S)=\begin{cases}\dfrac{y_S+1-\alpha}{\mu_F} & \text{if } y_S\le\mu_S\Big(\dfrac{1-\alpha}{\mu_F-\mu_S}\Big),\\[2ex] \dfrac{y_S}{\mu_S} & \text{otherwise.}\end{cases} \tag{B.5}$$

When $\mu_S=\mu_F$, we treat $\frac{1-\alpha}{\mu_F-\mu_S}$ as positive infinity, implying that $\mathcal{S}(\alpha;y_S)=\frac{y_S+1-\alpha}{\mu_F}$.

Proof. The fluid at position y_S in the buffer at station S will depart from station S after $T_S(y_S)=y_S/\mu_S$ units of time by equation (B.1). When this customer arrives at station F, the amount of fluid in the buffer there is $Q_F(\alpha;y_S/\mu_S)$, which we can determine using equation (B.3). We can therefore write her total system time $\mathcal{S}(\alpha;y_S)$ as

$$\begin{aligned} \mathcal{S}(\alpha;y_S) &= \frac{y_S}{\mu_S}+T_F\Big(Q_F\Big(\alpha;\frac{y_S}{\mu_S}\Big)\Big) \\ &= \frac{y_S}{\mu_S}+\frac{1}{\mu_F}\Big[1-\alpha+\min\{\alpha,y_S\}-\frac{\mu_F y_S}{\mu_S}\Big]^+ \\ &= \frac{y_S}{\mu_S}+\frac{1}{\mu_F}\Big[1-\alpha+y_S-\frac{\mu_F y_S}{\mu_S}\Big]^+, \end{aligned} \tag{B.6}$$

where the last substitution follows from the fact that for any SF customer we must have $y_S\le\alpha$. The bracketed term in equation (B.6) is nonnegative only if $y_S\le\mu_S(1-\alpha)/(\mu_F-\mu_S)$, in which case the inequality holds for any $\hat{y}_S$). Note that we define the RHS of the inequality as ∞ if $\mu_S=\mu_F$, for notational convenience:

the inequality always holds in that case, correctly implying that the first piece of equation (B.1) always governs the function $\mathcal{S}$ for equal service rates. Substitution then yields the first expression in equation (B.5). Otherwise, i.e., if $y_S > \mu_S(1-\alpha)/(\mu_F - \mu_S)$, taking the positive part of the bracketed term gives zero, and we are left with y_S/μ_S. We conclude equation (B.5). $\square$

The piecewise nature of $\mathcal{S}(\alpha; y_S)$ arises because the buffer at station F empties faster than that at station S. A customer near the front of the station S buffer will depart station S to find some customers still in the station F buffer, so her system time is determined by how long it takes station F to process both all of the FS customers and the SF customers in front of her. By contrast, a customer near the back of the station S buffer will arrive to station F after it clears and find it empty; for this infinitesimal customer, since her own service times are negligible, her system time is determined by how long station S takes to process the SF customers in front of her.

We similarly denote by $\mathcal{F}(\alpha; y_F)$ the total system time for the infinitesimal customer starting in position y_F in the station F queue, given a total SF fraction α. Because the station S buffer empties at a slower rate, all FS customers will be delayed in this buffer when they depart station F. The system time for these customers can thus be expressed more simply.

Lemma B.2 **(System Time for *FS* Customers)**. *We can express the function $\mathcal{F}(\alpha; y_F)$ by*

$$\mathcal{F}(\alpha; y_F) = \frac{\alpha + y_F}{\mu_S}. \tag{B.7}$$

Proof. The proof proceeds along similar lines to the proof of Lemma B.1. The fluid at position y_F in the buffer at station F will depart from station F after $T_F(y_F) = y_F/\mu_F$ units of time by equation (B.2). When this customer arrives at station S, the amount of fluid in the buffer there is $Q_S(\alpha; y_F/\mu_F)$, which we can determine using equation (B.4). The total system time for this customer is then

$$\begin{aligned}
\mathcal{F}(\alpha; y_F) &= \frac{y_F}{\mu_F} + T_S\Big(Q_S\Big(\alpha; \frac{y_F}{\mu_F}\Big)\Big) \\
&= \frac{y_F}{\mu_F} + \frac{1}{\mu_S}\Big[\alpha + \min\{1-\alpha, y_F\} - \frac{\mu_S y_F}{\mu_F}\Big]^+ \\
&= \frac{y_F}{\mu_F} + \frac{1}{\mu_S}\Big(\alpha + y_F - \frac{\mu_S y_F}{\mu_F}\Big) \\
&= \frac{\alpha + y_F}{\mu_S}.
\end{aligned}$$

The equivalence between the second and third equations holds because we must have $y_F \le 1-\alpha$, and because $\mu_S/\mu_F \le 1$ then implies that the bracketed term is nonnegative. $\square$

Lastly, we prove Property 1.

Proof of Property 1. When $\alpha > \mu_S/\mu_F$, we have that $1/\mu_F < \alpha/\mu_S$ which, combined with Lemma B.1 (with $y_S = \alpha$), implies that $\sup\{\text{supp}(S(\alpha)\} = \alpha/\mu_S$. Also, by Lemma B.2, taking $y_F = 0$, we get that $\inf\{\text{supp}(F(\alpha))\} = \alpha/\mu_S$. This proves equation (5).

By equation (5), note that for any $0 < k < K$, $\inf\{\text{supp}(S_i)\} = \inf\{\text{supp}(F_j)\} = \alpha/\mu_S$ for any $1 \le i \le k$, $1 \le j \le K-k$. Thus, we have that

$$\pi_k(\alpha) = \mathsf{Pr}\Big[\frac{\sum_{i=1}^k S_i(\alpha)}{k} \le \frac{\sum_{j=1}^{K-k} F_j(\alpha)}{K-k}\Big] = 1. \qquad \square$$

C. Proof of Claim 1

Proof. Suppose $\alpha = 1$, and consider an arbitrary customer contemplating deviation from route SF to route FS. For $\alpha = 1$ and $\mu_S < \mu_F$, the RHS of the condition for the first piece of equation (B.5) is zero, so we can focus on the second piece. Note that if $\mu_S = \mu_F$, although the first piece always governs the function, the expressions for the first and second piece are equal when $\alpha = 1$. Thus, we can use the second piece of equation (B.5) for our analysis for all $\mu_S \le \mu_F$. For a given starting position y_S in the station S buffer, the system time is then y_S/μ_S by Lemma B.1. Because customers in each buffer are sequenced uniformly at random, the expected system time for an SF customer given $\alpha = 1$ is equal to $1/(2\mu_S)$. If an arbitrary nonatomic customer deviates to route FS, then her position in the station F buffer will be $y_F = 0$. By Lemma B.2, her expected system time will thus be $1/\mu_S > 1/(2\mu_S)$, implying that she has no incentive to deviate to route FS. We conclude that $\pi^{FR}(1) = 1$.

Next, suppose $\alpha = 0$ and $2\mu_S \ge \mu_F$. First, we note that if $\mu_S = \mu_F$, then $\pi^{FR}(0) = 0$ by symmetry. For the rest of the proof, we assume that $\mu_S < \mu_F$. If a customer contemplating deviation remains on route FS, then by Lemma B.2 and the fact that her starting position is uniform in the station F buffer, her expected system time is $1/(2\mu_S)$. If she deviates to route SF, then her position in the station S buffer will be $y_S = 0$. By Lemma B.1, then, her system time will be $1/\mu_F \ge 1/(2\mu_S)$, where the inequality holds by our assumption that $2\mu_S \ge \mu_F$. Thus, the customer has no incentive to deviate from route FS, and we conclude that $\pi^{FR}(0) = 0$. □

D. Details for Cumulative System Time

Here, we study in more detail the cumulative system time as a function of the SF fraction α. We first derive this function in closed form, and then we study its value at α^* to reveal additional insights.

D.1. Expression for Cumulative System Time

The following proposition gives the expression for the cumulative system time. Importantly, the cumulative system time under herding is the same whether on route SF or on route FS.

Proposition D.1 **(Cumulative System Time)**. *Let*

$$\xi_\alpha := \min\left\{\alpha, \mu_S\left(\frac{1-\alpha}{\mu_F - \mu_S}\right)\right\}. \tag{D.1}$$

The cumulative system time is given by

$$D(\alpha, \mu_S, \mu_F) = \begin{cases} \dfrac{\alpha(1-\alpha)}{\mu_S} + \dfrac{(1-\alpha)^2}{2\mu_S} + \dfrac{\alpha(1-\alpha)}{\mu_F} + \dfrac{\alpha^2}{2\mu_F} & \text{if } \alpha \le \dfrac{\mu_S}{\mu_F}, \\ \dfrac{\alpha(1-\alpha)}{\mu_S} + \dfrac{(1-\alpha)^2}{2\mu_S} + \dfrac{\xi_\alpha(1-\alpha)}{\mu_F} + \dfrac{\alpha^2}{2\mu_S} + \xi_\alpha^2\left(\dfrac{1}{2\mu_F} - \dfrac{1}{2\mu_S}\right) & \text{otherwise.} \end{cases} \tag{D.2}$$

Proof. The overall cumulative system time can be broken down into two components: the cumulative system time for SF customers and the cumulative system time for FS customers. Respectively denoting these by $d_S(\cdot)$ and $d_F(\cdot)$, we have

$$D(\alpha, \mu_S, \mu_F) = d_S(\alpha, \mu_S, \mu_F) + d_F(\alpha, \mu_S, \mu_F). \tag{D.3}$$

First, we compute the cumulative system time for FS customers. By Lemma B.2, for a customer in position y_F, the system time is

$$\mathcal{F}(\alpha; y_F) = \frac{\alpha + y_F}{\mu_S}.$$

Integrating over all customers, we get that the cumulative system time for FS customers is

$$d_F(\alpha,\mu_S,\mu_F)=\int_0^{1-\alpha}\frac{\alpha+y_F}{\mu_S}\mathrm{d}y_F=\frac{\alpha(1-\alpha)}{\mu_S}+\frac{(1-\alpha)^2}{2\mu_S}.$$

For SF customers, by Lemma B.1, for a customer in position y_S, the system time is

$$\mathcal{S}(\alpha;y_S)=\begin{cases}\dfrac{y_S+1-\alpha}{\mu_F} & y_S\le\mu_S\Big(\dfrac{1-\alpha}{\mu_F-\mu_S}\Big),\\ \dfrac{y_S}{\mu_S} & \text{otherwise.}\end{cases}$$

Recall from equation (D.1) that

$$\xi_\alpha=\min\Big\{\alpha,\ \mu_S\Big(\frac{1-\alpha}{\mu_F-\mu_S}\Big)\Big\}.$$

Integrating, we get

$$\begin{aligned}d_S(\alpha,\mu_S,\mu_F)&=\int_0^{\alpha}\mathcal{S}(\alpha;y_S)\mathrm{d}y_S\\&=\int_0^{\xi_\alpha}\frac{y_S+1-\alpha}{\mu_F}\mathrm{d}y_S+\int_{\xi_\alpha}^{\alpha}\frac{y_S}{\mu_S}\mathrm{d}y_S\\&=\frac{\xi_\alpha(1-\alpha)}{\mu_F}+\frac{\xi_\alpha^2}{2\mu_F}+\frac{\alpha^2}{2\mu_S}-\frac{\xi_\alpha^2}{2\mu_S}.\end{aligned}$$

Note that if $\xi_\alpha=\alpha$, then the range of the second integral above is empty (coinciding with the latter two terms of the reduced expression canceling each other out), and we therefore conclude

$$d_S(\alpha,\mu_S,\mu_F)=\begin{cases}\dfrac{\alpha(1-\alpha)}{\mu_F}+\dfrac{\alpha^2}{2\mu_F} & \text{if }\alpha\le\dfrac{\mu_S}{\mu_F},\\ \dfrac{\xi_\alpha(1-\alpha)}{\mu_F}+\dfrac{\xi_\alpha^2}{2\mu_F}+\dfrac{\alpha^2}{2\mu_S}-\dfrac{\xi_\alpha^2}{2\mu_S} & \text{otherwise.}\end{cases}$$

Substituting our expressions for $d_S(\cdot)$ and $d_F(\cdot)$ into equation (D.3) gives the formula in equation (D.2). ☐

D.2. Numerical Study of Cumulative System Time

Next, we look at how the relative cumulative system time (compared to the baseline under herding with $\mu_S=\mu_F$) varies with the service rate ratio. Without loss of generality, we fix the sum of the service rates at 1. We study the cumulative system time as a function of x, where $\mu_S=x$ and $\mu_F=1-x$. We can think of the service provider as having a certain amount of capacity (servers, machines, etc.) to divide among the two stations. We will compare the system performance in equilibrium between the fully rational case and the procedurally rational case as the allocation x changes, for different values of the prior noise σ. It is convenient to relate x to the service rate ratio $r=\mu_S/\mu_F$. We have $r=x/(1-x)$, and therefore $x=r/(1+r)$. As x ranges from 0 to 1/2, the ratio r ranges from 0 to 1.

Figure D.1(a) illustrates how cumulative system time changes with service rate ratios under different values of σ (we fix $K=2$, and note that we plot only the interval $[.2,1]$ because the system time increases without bound as the ratio approaches zero). The impact of the standard deviation σ of the prior point estimates is somewhat predictable: for a given service ratio, the less noise in the prior, the closer the equilibrium is to herding, and thus the smaller the cumulative system time. The effect of changes in the service rate ratio for given σ is more nuanced. It is clear that, if customers will herd (the solid line in Figure D.1(a)), then the cumulative system time is minimized by setting the service rates exactly equal, i.e., $r=1$ and $\mu_S=\mu_F=1/2$. Interestingly, for procedurally rational equilibrium, the cumulative system time is non-monotonic in the service rate ratio. For low service rate ratios, the probability of an SF draw being worse than an FS draw

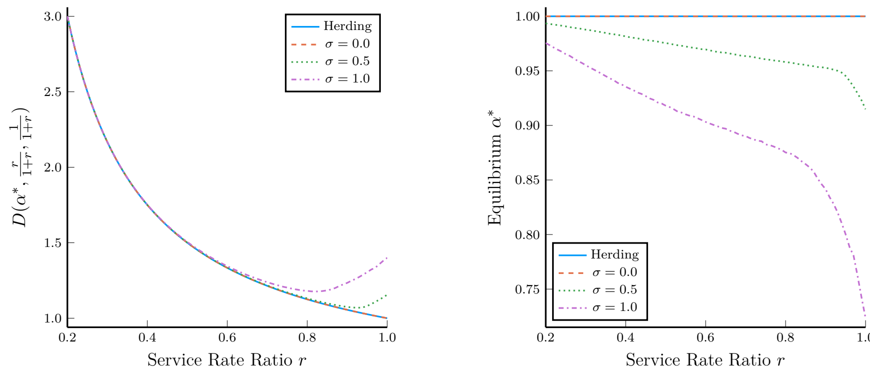

(a) Cumulative system time at α^* (b) Largest equilibrium α^*

Figure D.1 Cumulative system time and equilibria under herding and procedural rationality, versus $r = \mu_S/\mu_F$.

is small for almost any α, so the equilibrium under procedural rationality is very close to herding (see Figure D.1(b)). Thus, as the service rate ratio increases, the benefit of more balanced service rates causes the cumulative system time at equilibrium to improve. On the other hand, when the service rates get closer together, the advantage of route SF diminishes as the supports of the system times overlap more (and the region $(\mu_S/\mu_F, 1)$, where they intersect only at a single point, shrinks). Hence, for given α, the probability of an SF draw being worse than an FS draw increases, so the equilibrium α^* moves away from herding. Eventually, the impact of the deviation from herding outweighs the benefit of more balanced service rates, and at this point the cumulative system time at equilibrium starts to increase with the service rate ratio. This implies that if we let the firm pick the ratio, under procedural rationality, the optimal ratio is strictly below 1, unlike with fully rational customers. Encouragingly, the attendant performance loss diminishes as the noise in the prior decreases. From Figure D.1(a), for $\sigma = 1$, the cumulative system time at the optimal ratio is 17.7% worse than that at the optimal ratio under herding; by contrast, for $\sigma = 0.5$ the degradation is only 7.0%, and for $\sigma = 0$, there is no performance loss.

E. A Concrete Case of General Route Anecdotes: Sampling with Discernibility

The general route anecdotes sampling process discussed in the paper and analyzed in Section 4.2 allows the numbers K_S and K_F of SF and FS anecdotes to follow an arbitrary joint distribution. Proposition 4 shows that in this general model, as long as the probability of obtaining both types of anecdotes approaches 1, the equilibrium α^* will converge to herding. The general route anecdotes sampling process is somewhat abstract, so here we provide a specific example of it for concreteness.

Let us return to the social media example described in the introduction, in which a customer requests information from other customers about the system times of different routes. After her first post, suppose that the only route reflected in the replies is the route visiting A first and then B. In this case, the guest may post again, asking specifically "did anyone go to B first and then A? What were your wait times?" The activeness and size of the group and the willingness of the users to help others both moderate the probability that her second post (about the missing route) will receive replies. For a given population proportion of routes, she will be more likely to receive a reply about the missing route if the group is larger and more active.

To incorporate sampling that is biased to some degree toward the missing type of anecdote, we consider a sequential sampling model that we call "sampling with discernibility." Under sampling with discernibility, the process for customers to obtain their anecdotes is as follows. Each customer will draw a sample of anecdotes

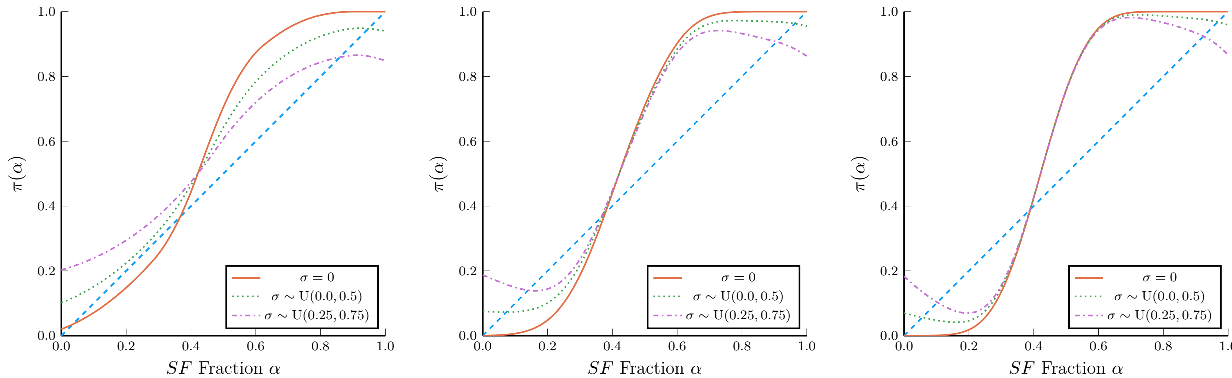

(a) $K = 2$ (b) $K = 5$ (c) $K = 8$

Figure F.1 Procedurally rational response function $\pi(\alpha)$ for random σ and different K ($\mu_S = .9$, $\mu_F = 1$).

of size K with replacement from the rest of the customers. The first anecdote a customer draws is selected uniformly at random, and thus, the probability of drawing a type SF anecdote first is α. On any draw after the first, if one type of anecdotes is missing, then the next anecdote is biased towards the missing type. The bias is modeled with a *discernibility* parameter $\beta \geq 1$. This parameter can be viewed as a simplified abstraction of the factors other than the population proportion of routes (like the size and activeness of the online group discussed above) that influence the probability of obtaining the missing type of anecdote. We then let the probability that a type FS anecdote is drawn given that the current sample does not contain type FS be

$$\frac{\beta(1-\alpha)}{\alpha + \beta(1-\alpha)}. \tag{E.1}$$

Similarly, given that the current sample does not contain type SF, we let the probability that a type SF anecdote is drawn be

$$\frac{\beta\alpha}{\beta\alpha + (1-\alpha)}. \tag{E.2}$$

After obtaining both types of anecdotes, a customer reverts to sampling uniformly from the population. When $\beta = 1$, customers do not sample in a way that is biased towards the missing anecdote, and we recover the random route anecdotes model. On the other hand, as $\beta \to \infty$, we approach perfect discernibility, as the probability of obtaining the anecdote of the missing type converges to 1 in the second draw. Applying the characteristic function argument in the proof of Lemma G.1, we can show that the response function under this sampling model is continuous. Thus, for any $K \geq 2$, Proposition 4 implies that α^* converges to 1 as discernibility increases.

F. Numerical Study of Customer Heterogeneity

Here, we study systems with heterogeneity among customers in the sample size and the prior noise. First, we consider a procedurally rational model with heterogeneity in the degree of noise in the prior: we suppose that σ is a uniformly distributed random variable. Figure F.1 plots the response function for random σ centered on the deterministic values in Figure 1, for the same sample sizes and service rates. The curves in Figure F.1 are extremely similar to—in fact, nearly indistinguishable from—the corresponding curves in Figure 1. This indicates, encouragingly, that randomness in σ does not change the qualitative nature of the procedurally rational response function.

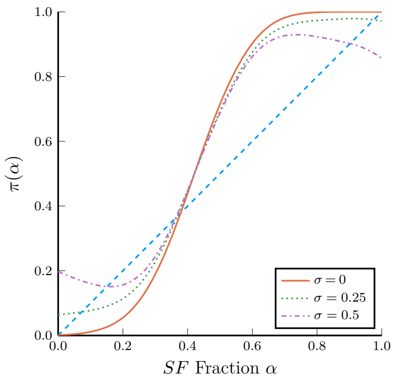

(a) $K \sim \mathrm{U}\{3,7\}$

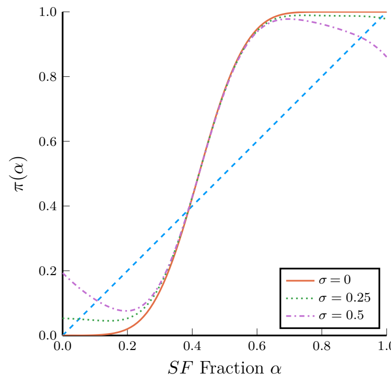

(b) $K \sim \mathrm{U}\{6,10\}$

Figure F.2 Procedurally rational response function $\pi(\alpha)$ for different σ and random K ($\mu_S = .9$, $\mu_F = 1$). For integers $x < y$, U$\{x,y\}$ denotes a discrete uniform distribution on the integers $x, x+1, \ldots, y-1, y$.

We also consider a procedurally rational model with a random sample size. In this model, for each infinitesimal customer, the sample size is a discrete uniform random variable. Figure F.2 plots the response function for random sample sizes centered on the deterministic values considered in the right two panels of Figure 1, for the same service rates. As was the case with random σ, for random sample sizes the curves in Figure F.2 are virtually indistinguishable from the corresponding curves in Figure 1. The only discernible difference is that randomness in the sample size brings the response function (very) slightly further from the extremes of 0 and 1, as expected with more randomness.

G. Proof of Proposition 1

The proof requires some auxiliary lemmas. We first show with Lemma G.1 that the response function π is continuous in α. We then prove with Lemma G.2 that the set of procedurally rational equilibria is non-empty and compact. Next, Lemma G.3 establishes a useful result for bounding the largest equilibrium. Most of what remains to prove Proposition 1 is then accomplished by Lemma G.4, which shows that for $\alpha \in (1/2, 1)$, as K increases, eventually both (i) customers will receive many anecdotes from both routes and (ii) their corresponding estimates of the system times for each route will be accurate enough that they will choose route SF with high probability. From there, the proof of Proposition 1 is straightforward. Lastly, we also prove Corollary 1, that for any $\alpha \in (1/2, 1)$, the procedurally rational response function π converges to the fully rational response function π^{FR} (which prescribes herding, i.e., $\alpha = 1$) as $K \to \infty$.

Lemma G.1 **(Continuity of π).** *For any fixed K, σ, the function $\pi(\alpha)$ is continuous over* $[0,1]$.

Proof. By definition, the distributions of S_j and F_j depend on α for all $0 \le j \le K$. To highlight this dependence, in this proof, we use $S_j(\alpha)$ and $F_j(\alpha)$ in place of S_j and F_j. Let

$$\pi_0(\alpha) = \mathsf{Pr}\left[S_0(\alpha) \le \frac{\sum_{j=1}^K F_j(\alpha)}{K}\right], \pi_K(\alpha) = \mathsf{Pr}\left[\frac{\sum_{j=1}^K S_j(\alpha)}{K} \le F_0(\alpha)\right],$$

$$\text{and } \pi_k(\alpha) = \mathsf{Pr}\left[\frac{\sum_{i=1}^k S_i(\alpha)}{k} \le \frac{\sum_{j=1}^{K-k} F_j(\alpha)}{K-k}\right], \forall k = 1, \ldots K-1.$$

Recall from (2) that

$$\pi(\alpha) = (1-\alpha)^K \pi_0(\alpha) + \alpha^K \pi_K(\alpha) + \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S} (1-\alpha)^{K-k_S} \pi_{k_S}(\alpha).$$

Thus, it is sufficient to prove that $\pi_k(\alpha)$ is continuous in α over $[0,1]$ for any $k = 0, \dots, K$.

To prove the continuity of $\pi_k(\alpha)$, observe that for any sequence $\{\alpha_i\}_{i=1}^{\infty}$ where $\lim_{i\to\infty} \alpha_i = \alpha$, we have that $\lim_{i\to\infty} S_k(\alpha_i)$ (and $\lim_{i\to\infty} F_k(\alpha_i)$) converges in distribution to $S_k(\alpha)$ (and $F_k(\alpha)$) for each $0 \le k \le K$. Therefore, by Lévy's continuity theorem, we have that the characteristic function of $S_k(\alpha_i)$ (and $F_k(\alpha_i)$) converges to the characteristic function $S_k(\alpha)$ (and $F_k(\alpha)$).

Observe that for $\pi_0(\alpha)$, we have

$$\pi_0(\alpha) = \mathsf{Pr}[S_0(\alpha) - K^{-1} \sum_{j=1}^{K} F_j(\alpha) \le 0].$$

for any sequence $\{\alpha_i\}_{i=1}^{\infty}$ where $\lim_{i\to\infty} \alpha_i = \alpha$. Because $S_k(\alpha)$ and $F_k(\alpha)$ for $0 \le k \le K$ are independent, the characteristic function of $S_0(\alpha_i) - K^{-1} \sum_{j=1}^{K} F_j(\alpha_i)$ is a product of characteristic functions of $S_0(\alpha_i)$ and $\{F_j(\alpha_i)\}_{1\le j\le K}$. Hence, it must converge to the characteristic function of $S_0(\alpha) - K^{-1} \sum_{j=1}^{K} F_j(\alpha)$. This implies that

$$\lim_{i\to\infty} \pi_0(\alpha_i) = \lim_{i\to\infty} \mathsf{Pr}[S_0(\alpha_i) - K^{-1} \sum_{j=1}^{K} F_j(\alpha_i) \le 0] = \mathsf{Pr}[S_0(\alpha) - K^{-1} \sum_{j=1}^{K} F_j(\alpha) \le 0] = \lim_{i\to\infty} \pi_0(\alpha).$$

Note that the continuity of $\pi_0(\alpha)$ relied on the property that $\pi_0(\alpha)$ can be expressed as a probability of a weighted, independent sum of $S_k(\alpha)$ and $F_k(\alpha)$ being less than or equal to zero. Similarly, observe that for each $k = 1, \dots, K$, the function $\pi_k(\alpha)$ can also be expressed as the probability that a weighted, independent sum of $S_k(\alpha)$ and $F_k(\alpha)$ is less than or equal to zero. Therefore, applying a similar argument, for any sequence where $\lim_{i\to\infty} \alpha_i = \alpha$, we have

$$\lim_{i\to\infty} \pi_k(\alpha_i) = \pi_k(\alpha), \forall k = 1, \dots, K.$$

□

Lemma G.2 **(Set of Equilibria Is Non-Empty and Compact)**. *For any K and $\sigma \ge 0$, the set of procedurally rational equilibria, i.e., $\{\alpha : \pi(\alpha) = \alpha\}$, is non-empty and compact.*

Proof. Note that $\pi(\alpha)$ maps the compact set $[0,1]$ to $[0,1]$ and is continuous by Lemma G.1. By Brouwer's fixed-point theorem, we have that there exists at least one $0 \le \alpha^* \le 1$ such that $\pi(\alpha^*) = \alpha^*$.

Also, by Lemma G.1, $\pi(\alpha)$ is continuous and hence $h(\alpha) := \pi(\alpha) - \alpha$ is continuous over $[0,1]$. The set of procedurally rational equilibria is the preimage $h^{-1}(\{0\})$. Because $h(\alpha)$ is continuous over $[0,1]$, the set of equilibria is therefore closed because the preimage of a closed set (here the singleton $\{0\}$) under a continuous function is closed. The set of equilibria is also bounded as it is contained in $[0,1]$, and it is therefore compact.

□

Lemma G.3. *Suppose there is some $0 < \alpha < 1$ such that $\pi(\alpha) > \alpha$. Then we must have $\alpha^* > \alpha$.*

Proof. If $\pi(1) = 1$, then $\alpha^* = 1$ and the result holds. If instead $\pi(1) < 1$, then because π is continuous and $\pi(\alpha) > \alpha$, the intermediate value theorem implies that there exists α' with $\alpha < \alpha' < 1$ such that $\pi(\alpha') = \alpha'$, implying that $\alpha^* \ge \alpha' > \alpha$. □

LEMMA G.4 **(Limit of Probabilities in K).** *Consider a fixed $\alpha \in (\frac{1}{2}, 1)$. We have*

$$\lim_{K\to\infty} \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \pi_{k_S}(\alpha) = 1.$$

Proof. Let $\bar{s}_{K|k_S}(\alpha) = \sum_{i=1}^{k_S} S_i(\alpha)/k_S$ and $\bar{f}_{K|k_S}(\alpha) = \sum_{j=1}^{K-k_S} F_j(\alpha)/(K-k_S)$ be the conditional average waiting time for the SF and FS anecdotes, respectively, conditioned on drawing exactly k_S SF anecdotes (and thus exactly $K-k_S$ FS anecdotes). From equation (3), we can then write $\pi_{k_S}(\alpha) = \mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) \le \bar{f}_{K|k_S}(\alpha)]$. Observe that

$$\begin{aligned}
\sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \pi_{k_S}(\alpha) &= \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) \le \bar{f}_{K|k_S}(\alpha)] \\
&= \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} (1-\mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) > \bar{f}_{K|k_S}(\alpha)]) \\
&= 1-(1-\alpha)^K - \alpha^K - \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) > \bar{f}_{K|k_S}(\alpha)].
\end{aligned}$$

Since $1/2 < \alpha < 1$, we have $\lim_{K\to\infty}(1-\alpha)^K = \lim_{K\to\infty} \alpha^K = 0$. Therefore, we can prove the lemma by showing that

$$\lim_{K\to\infty} \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) > \bar{f}_{K|k_S}(\alpha)] = 0.$$

For any fixed k_S, the binomial probability $\binom{K}{k_S}\alpha^{k_S}(1-\alpha)^{K-k_S}$ goes to 0 as K goes to infinity. So, we have that for any N_0 and any $\epsilon > 0$, there exists $K(N_0)$ such that for all $K \ge K(N_0)$, we have

$$\sum_{k_S=1}^{N_0} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} + \sum_{k_S=K-N_0}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} < \frac{\epsilon}{3}. \tag{G.1}$$

Let s^* (f^*) be the true expected system time for route SF (FS), given the fixed SF fraction $\alpha \in (1/2, 1)$. By Lemma B.1, we have

$$s^* = \frac{1}{\alpha}\int_0^\alpha \mathcal{S}(\alpha; y_S)\mathrm{d}y_S \le \frac{1}{\alpha}\int_0^\alpha \frac{y_S + 1 - \alpha}{\mu_S}\mathrm{d}y_S = \frac{1-\alpha/2}{\mu_S}; \tag{G.2}$$

and by Lemma B.2, the expected FS system time f^* is equal to

$$f^* = \frac{1}{1-\alpha}\int_0^{1-\alpha} \mathcal{F}(\alpha; y_F)\mathrm{d}y_F = \frac{1}{1-\alpha}\int_0^{1-\alpha} \frac{\alpha + y_F}{\mu_S}\mathrm{d}y_F = \frac{\alpha/2 + 1/2}{\mu_S}. \tag{G.3}$$

Note that $\alpha \in (1/2, 1)$ implies that $1-\alpha/2 < \alpha/2 + 1/2$. Combining this with equations (G.2) and (G.3), we have that $d := f^* - s^* > 0$.

By the weak law of large numbers (see, e.g., Grimmett and Stirzaker 2020, Section 7.4), for any $\epsilon > 0$, there exists N_0 such that for all K and k_S with $N_0 \le k_S \le K - N_0$, we have that

$$\mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) \ge s^* + \frac{d}{2}] < \frac{\epsilon}{3}, \text{ and } \mathsf{Pr}[\bar{f}_{K|k_S}(\alpha) \le f^* - \frac{d}{2}] < \frac{\epsilon}{3},$$

which in turn implies that

$$\mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) > \bar{f}_{K|k_S}(\alpha)] < \frac{2\epsilon}{3}. \tag{G.4}$$

Now, for any $\epsilon > 0$, choose N_0 so that inequality (G.4) is satisfied for all K and k_S with $N_0 \le k_S \le K - N_0$, and pick $K_0 = K(N_0)$ such that inequality (G.1) holds for all $K \ge K_0$. Then, for $K \ge K_0$, we have

$$\begin{aligned}
\sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) > \bar{f}_{K|k_S}(\alpha)] &\le \sum_{k_S=1}^{N_0} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} + \sum_{k_S=K-N_0}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \\
&\quad + \sum_{k_S=N_0+1}^{K-N_0-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) > \bar{f}_{K|k_S}(\alpha)] \\
\text{By (G.1),} \qquad &< \frac{\epsilon}{3} + \sum_{k_S=N_0+1}^{K-N_0-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) > \bar{f}_{K|k_S}(\alpha)] \\
\text{By (G.4),} \qquad &\le \frac{\epsilon}{3} + \frac{2\epsilon}{3} \sum_{k_S=N_0+1}^{K-N_0-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \\
&< \epsilon.
\end{aligned}$$

This establishes

$$\lim_{K\to\infty} \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \mathsf{Pr}[\bar{s}_{K|k_S}(\alpha) > \bar{f}_{K|k_S}(\alpha)] = 0,$$

which completes the proof of the lemma. □

Proof of Proposition 1. By Lemma G.4, for any $\alpha \in (1/2, 1)$, there exists K_0 such that for $K \ge K_0$,

$$\sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \pi_{k_S}(\alpha) > \alpha,$$

which implies that $\pi(\alpha) > \alpha$ by equation (4). Lemma G.3 then implies that $\alpha^* > \alpha$. Because this is true for any $1/2 \le \alpha < 1$, it implies that $\liminf_{K\to\infty} \alpha^* \ge 1$. Also, by definition, the largest equilibrium is no larger than 1, thus $\limsup_{K\to\infty} \alpha^* \le 1$. Thus, $\lim_{K\to\infty} \alpha^* = 1$. □

Proof of Corollary 1. By equations (G.2) and (G.3) from the proof of Lemma G.4, we have that for $\alpha \in (1/2, 1)$ the expected system time of SF is strictly less than the expected system time of FS. Thus, $\pi^{FR}(\alpha) = 1$ for $\alpha \in (1/2, 1)$.

Next, for $\alpha \in (1/2, 1)$, $\pi(\alpha) \le 1$ for any K, and by Lemma G.4 and equation (4), we have

$$\lim_{K\to\infty} \pi(\alpha) \ge \lim_{K\to\infty} \sum_{k_S=1}^{K-1} \binom{K}{k_S} \alpha^{k_S}(1-\alpha)^{K-k_S} \pi_{k_S}(\alpha) = 1.$$

Therefore, we have that for $\alpha \in (1/2, 1)$, $\lim_{K\to\infty} \pi(\alpha) = \pi^{FR}(\alpha) = 1$. □

H. Proof of Proposition 2

For any $\alpha \ge \mu_S/\mu_F$, by Equation (4) and Property 1, we have

$$\pi(\alpha) = \alpha^K \pi_0(\alpha) + (1-\alpha)^K \pi_K(\alpha) + (1 - \alpha^K - (1-\alpha)^K). \tag{H.1}$$

Let $\hat{S}_K(\alpha)$ be the estimated system time for route SF when all K anecdotes are SF type, and $\hat{F}_K(\alpha)$ be the estimated system time for route FS when all K anecdotes are FS type. Applying Property 1, we obtain that

$$\mathsf{E}[S_0] < \sup\{\text{supp}(\hat{S}_K(\alpha))\} = \mathcal{S}(\alpha;\alpha) = \alpha/\mu_S = \inf\{\text{supp}(\hat{F}_K(\alpha))\} < \mathsf{E}[F_0].$$

Therefore, as S_0 and F_0 are normally distributed, we have that

$$\pi_0(\alpha) = \mathsf{Pr}\left[S_0 \le \frac{\sum_{j=1}^K F_j}{K}\right] = \mathsf{Pr}\left[S_0 \le \hat{F}_K(\alpha)\right] > \mathsf{Pr}\left[S_0 \le \mathsf{E}[S_0]\right] = 1/2,$$

$$\pi_K(\alpha) = \mathsf{Pr}\left[\frac{\sum_{j=1}^K S_j}{K} \le F_0\right] = \mathsf{Pr}\left[\hat{S}_K(\alpha) \le F_0\right] > \mathsf{Pr}\left[\mathsf{E}[F_0] \le F_0\right] = 1/2.$$

Therefore, we have that

$$\pi(\alpha) > \frac{1}{2}(\alpha^K + (1-\alpha)^K) + (1 - \alpha^K - (1-\alpha)^K) = 1 - \frac{1}{2}\alpha^K - \frac{1}{2}(1-\alpha)^K. \tag{H.2}$$

To prove the first part of the proposition, by (H.2), when $\mu_S/\mu_F < 1/\sqrt{2} = \alpha$, we have

$$\pi(\alpha) > 1 - \frac{1}{2}\alpha^K - \frac{1}{2}(1-\alpha)^K \ge 1 - \frac{1}{2}\alpha^2 - \frac{1}{2}(1-\alpha)^2 = \frac{1}{2} + \alpha - \alpha^2 = 1/\sqrt{2}.$$

By Lemma G.3, we thus have that $\alpha^* > \alpha = 1/\sqrt{2}$.

We next prove the second part of the proposition. Observe that $K \ge (\frac{\mu_F}{\mu_F - \mu_S})^2$ is equivalent to $1 - \frac{1}{\sqrt{K}} \ge \mu_S/\mu_F$. Therefore, if $\alpha = 1 - \frac{1}{\sqrt{K}}$, we have $\alpha \ge \mu_S/\mu_F$. By (H.2), we have

$$\pi(\alpha) > 1 - \frac{1}{2}\alpha^K - \frac{1}{2}(1-\alpha)^K,$$

which implies

$$\pi(\alpha) - \alpha > 1 - \alpha - \frac{1}{2}\alpha^K - \frac{1}{2}(1-\alpha)^K =: f(\alpha). \tag{H.3}$$

Now, define $\tau_K := 1 - 1/\sqrt{K}$. We have

$$\tau_K^K = \left(1 - \frac{1}{\sqrt{K}}\right)^K = \left(1 + \frac{(-\sqrt{K})}{K}\right)^K \le \frac{1}{e^{\sqrt{K}}} < \frac{1}{\sqrt{K}},$$

where the first inequality follows from a basic property of the exponential function (see, e.g., Wang 1989) and the second inequality follows from the power series expansion for the exponential function; the expansion gives $e^{\sqrt{K}} = 1 + \sqrt{K} + K/2! + K^{3/2}/3! + \cdots > \sqrt{K}$, which implies $1/e^{\sqrt{K}} < 1/\sqrt{K}$. Noting also that $K \ge 2$ implies $1/(\sqrt{K})^K < 1/\sqrt{K}$, we conclude

$$f(\tau_K) = 1 - \tau_K - \frac{(1-\tau_K)^K}{2} - \frac{\tau_K^K}{2} = \frac{1}{\sqrt{K}} - \frac{1}{2}\left(\frac{1}{(\sqrt{K})^K} + \tau_K^K\right) > 0, \tag{H.4}$$

which then implies $\pi(\tau_K) > \tau_K$ by equation (H.3). Lemma G.3 then implies that $\alpha^* > \tau_K = 1 - \frac{1}{\sqrt{K}}$, completing the proof. □

I. Proof of Proposition 3

For any $\alpha \ge \mu_S/\mu_F$, by Property 1, we have $\pi_k(\alpha) = 1$ for any $1 \le k \le K-1$. Thus, we can rewrite equation (4) as

$$\pi(\alpha) = (1-\alpha)^K \pi_0(\alpha) + \alpha^K \pi_K(\alpha) + (1 - \alpha^K - (1-\alpha)^K),$$

where

$$\pi_0(\alpha) = \mathsf{Pr}\left[S_0 \le \frac{\sum_{j=1}^K F_j}{K}\right], \ \pi_K(\alpha) = \mathsf{Pr}\left[F_0 \ge \frac{\sum_{j=1}^K S_j}{K}\right].$$

Let $\hat{S}_K(\alpha) = \frac{\sum_{j=1}^K S_j}{K}$ and $\hat{F}_K(\alpha) = \frac{\sum_{j=1}^K F_j}{K}$. Recall that S_j had distribution $S(\alpha)$ and F_j had distribution $F(\alpha)$ for each $1 \le j \le K$. As $\mathsf{E}[S_0]$ and $\mathsf{E}[F_0]$ are equal to the expected values of $S(\alpha)$ and $F(\alpha)$, by Property 1, we have that

$$\mathsf{E}[S_0] < \sup\{\mathrm{supp}(\hat{S}_K(\alpha)\} = \mathcal{S}(\alpha;\alpha) = \alpha/\mu_S = \inf\{\mathrm{supp}(\hat{F}_K(\alpha))\} < \mathsf{E}[F_0].$$

Recall that S_0 is normally distributed with standard deviation σ. Letting Z be a standard normal random variable, we have

$$\lim_{\sigma\to 0} \pi_0(\alpha) = \lim_{\sigma\to 0} \mathsf{Pr}\left[\sigma Z \le \frac{\sum_{j=1}^K F_j}{K} - \mathsf{E}[S_0]\right] = 1,$$

where the last equality follows from the fact that $\mathsf{E}[S_0] < \inf\{\mathrm{supp}(\hat{F}(\alpha))\}$. Similarly, we also have that F_0 is normally distributed with standard deviation σ, and thus

$$\lim_{\sigma\to 0} \pi_K(\alpha) = \lim_{\sigma\to 0} \mathsf{Pr}\left[\frac{\sum_{j=1}^K S_j}{K} - \mathsf{E}[F_0] \le \sigma Z\right] = 1.$$

Putting everything together, we have for any $\alpha \ge \mu_S/\mu_F$ that

$$\lim_{\sigma\to 0} \pi(\alpha) = 1.$$

Thus, for any $\epsilon > 0$ there exists σ_0 such that for all $\sigma < \sigma_0$, we have $\pi(\alpha) > \alpha$. By Lemma G.3, this implies that an equilibrium exists that is strictly larger than α, and because this holds for any $\alpha \ge \mu_S/\mu_F$, it implies that $\liminf_{\sigma\to 0} \alpha^* \ge 1$. Also, by definition, the largest equilibrium is no larger than 1, thus $\limsup_{\sigma\to 0} \alpha^* \le 1$, implying that $\lim_{\sigma\to 0} \alpha^* = 1$. □

J. Proof of Proposition 4

For the general sampling process parameterized by some β, the number of anecdotes from each route can have an arbitrary distribution, and their sum need not be a fixed constant. Similar to the convenience notation from the randomly sampled anecdotes case, for the general sampling process we define

$$\begin{aligned}
\pi_{0,0}(\alpha) &= \mathsf{Pr}\left[S_0(\alpha) \le F_0(\alpha)\right], \\
\pi_{0,k_F}(\alpha) &= \mathsf{Pr}\left[S_0(\alpha) \le \frac{\sum_{j=1}^{k_F} F_j(\alpha)}{k_F}\right], \forall k_F = 1,2,\ldots, \\
\pi_{k_S,0}(\alpha) &= \mathsf{Pr}\left[\frac{\sum_{j=1}^{k_S} S_j(\alpha)}{k_S} \le F_0(\alpha)\right], \forall k_S = 1,2,\ldots \\
\text{and } \pi_{k_S,k_F}(\alpha) &= \mathsf{Pr}\left[\frac{\sum_{i=1}^{k_S} S_i(\alpha)}{k_S} \le \frac{\sum_{j=1}^{k_F} F_j(\alpha)}{k_F}\right], \forall k_S = 1,2,\ldots, k_F = 1,2,\ldots.
\end{aligned}$$

Let $g_{\beta,k_S,k_F}(\alpha) := \mathsf{Pr}[K_S(\alpha) = k_S, K_F(\alpha) = k_F]$. The generic form for the response function, denoted by π^β, is then

$$\pi^\beta(\alpha) = \sum_{k_S=0}^\infty \sum_{k_F=0}^\infty g_{\beta,k_S,k_F}(\alpha)\pi_{k_S,k_F}(\alpha). \tag{J.1}$$

Consider some $\alpha \in (\mu_S/\mu_F, 1)$. Property 1 implies that $\sup\{\mathrm{supp}(S_i(\alpha)\} = \alpha/\mu_S$ for each of the S_i's and that $\inf\{\mathrm{supp}(F_j(\alpha))\} = \alpha/\mu_S$ for each of the F_j's. Similar to equation (6) from the random route anecdotes case, for general route anecdotes this implies

$$\pi_{k_S,k_F}(\alpha) = 1, \forall k_S = 1,2,\ldots, k_F = 1,2,\ldots,$$

and thus equation (J.1) becomes

$$\pi^{\beta}(\alpha)=\sum_{k_S=1}^{\infty}\sum_{k_F=1}^{\infty}g_{\beta,k_S,k_F}(\alpha)+\sum_{k_F=1}^{\infty}g_{\beta,0,k_F}(\alpha)\pi_{0,k_F}(\alpha)+\sum_{k_S=1}^{\infty}g_{\beta,k_S,0}(\alpha)\pi_{k_S,0}(\alpha)+g_{\beta,0,0}(\alpha)\pi_{0,0}(\alpha).$$

By definition, we have $\gamma(\alpha,\beta)=\sum_{k_S=1}^{\infty}\sum_{k_F=1}^{\infty}g_{\beta,k_S,k_F}(\alpha)$, which then implies

$$\pi^{\beta}(\alpha)=\gamma(\alpha,\beta)+\sum_{k_F=1}^{\infty}g_{\beta,0,k_F}(\alpha)\pi_{0,k_F}(\alpha)+\sum_{k_S=1}^{\infty}g_{\beta,k_S,0}(\alpha)\pi_{k_S,0}(\alpha)+g_{\beta,0,0}(\alpha)\pi_{0,0}(\alpha). \tag{J.2}$$

Because $\sum_{k_S=0}^{\infty}\sum_{k_F=0}^{\infty}g_{\beta,k_S,k_F}(\alpha)=1$, we have $\sum_{k_S=1}^{\infty}\sum_{k_F=1}^{\infty}g_{\beta,k_S,k_F}(\alpha)=1-(\sum_{k_S=1}^{\infty}g_{\beta,k_S,0}(\alpha)+\sum_{k_F=1}^{\infty}g_{\beta,0,k_F}(\alpha)+g_{\beta,0,0}(\alpha))$. Therefore, since $\lim_{\beta\to\infty}\sum_{k_S=1}^{\infty}\sum_{k_F=1}^{\infty}g_{\beta,k_S,k_F}(\alpha)=\lim_{\beta\to\infty}\gamma(\alpha,\beta)=1$, we must have $\lim_{\beta\to\infty}g_{\beta,0,k_F}(\alpha)=0$ and $\lim_{\beta\to\infty}g_{\beta,k_S,0}(\alpha)=0$ for any nonnegative integers k_S and k_F. Together with equation (J.2), this implies $\lim_{\beta\to\infty}\pi^{\beta}(\alpha)=\lim_{\beta\to\infty}\gamma(\alpha,\beta)=1$, and note that this holds for arbitrarily chosen $\alpha\in(\mu_S/\mu_F,1)$.

Since $\lim_{\beta\to\infty}\pi^{\beta}(\alpha)=1$ for $\alpha\in(\mu_S/\mu_F,1)$, for any such α there exists M_α such that for all $\beta>M_\alpha$, we have $\pi^{\beta}(\alpha)>\alpha$. Then, as π^{β} is continuous, an analogous version of Lemma G.3 for general route anecdotes implies that there exists an equilibrium that is strictly larger than α, and an analogous version of Lemma G.2 implies the existence of a largest equilibrium α^*. Because this is true for any $\mu_S/\mu_F<\alpha<1$, it implies that $\liminf_{\beta\to\infty}\alpha^*\geq 1$. Also, by definition, the largest equilibrium is no larger than 1, thus $\limsup_{\beta\to\infty}\alpha^*\leq 1$, implying that $\lim_{\beta\to\infty}\alpha^*=1$. □

K. Proof of Proposition 5

First, note that D is continuous in α because the two pieces of equation (D.2) coincide at the boundary. Differentiating the first piece of equation (D.2) twice gives

$$\frac{\partial^2 D}{\partial\alpha^2}=-\frac{1}{\mu_S}-\frac{1}{\mu_F}<0. \tag{K.1}$$

So, the function is concave in α on $[0,\mu_S/\mu_F]$ and its minimum over this interval will occur at one of the endpoints. We have $D(0,\mu_S,\mu_F)=1/(2\mu_S)$, while

$$D(\mu_S/\mu_F,\mu_S,\mu_F)=\frac{1}{2\mu_S}+\frac{\mu_S}{2\mu_F^2}\Big(1-\frac{\mu_S}{\mu_F}\Big)\geq\frac{1}{2\mu_S}. \tag{K.2}$$

For the case with $\mu_S=\mu_F$, we never have $\alpha>\mu_S/\mu_F$, and the minimizer on $[0,\mu_S/\mu_F]$ is the minimizer on $[0,1]$. Evaluating the second endpoint, we have $D(1,\mu_S,\mu_F)=1/(2\mu_S)=D(0,\mu_S,\mu_F)$. So, herding on either route minimizes the cumulative system time over $[0,1]$.

For the rest of the proof, we assume that $\mu_S<\mu_F$. In this case, on $[0,\mu_S/\mu_F]$, the function has a unique minimizer of 0 by equation (K.2), with a cumulative system time of $1/(2\mu_S)$. We must also consider the second piece of the function that governs on $[\mu_S/\mu_F,1]$. Differentiating the second piece of equation (D.2), we have

$$\frac{\partial D}{\partial\alpha}=-\frac{\mu_S(1-\alpha)}{\mu_F(\mu_F-\mu_S)}\leq 0 \text{ for } \alpha\leq 1. \tag{K.3}$$

Therefore, the cumulative system time is decreasing for $\alpha\in[\mu_S/\mu_F,1]$, and it is thus minimized at $\alpha=1$ on this interval. Substituting $\alpha=1$ into the second piece of the function gives $D(1,\mu_S,\mu_F)=1/(2\mu_S)=D(0,\mu_S,\mu_F)$. Thus, the cumulative system time is again minimized at $\alpha=0$ or 1, i.e., when customers herd on the same route. □

L. Prior Centered Elsewhere

For a given fixed route (either SF or FS), suppose that the prior is centered anywhere in $\text{supp}(\hat{S}_K(\alpha))$ for route SF and anywhere in $\text{supp}(\hat{F}_K(\alpha))$ for route FS (so not necessarily at the true expected system time)

Now, we observe that the proofs of Propositions 1, 4, D.1, and 5 and those of their supporting results do not rely on any properties of the prior, while the proofs of Propositions 2 and 3 use only the consequence of Property 1 that

$$\mathsf{E}[S_0] < \sup\{\text{supp}(\hat{S}_K(\alpha))\} = \mathcal{S}(\alpha;\alpha) = \alpha/\mu_S = \inf\{\text{supp}(\hat{F}_K(\alpha))\} < \mathsf{E}[F_0],$$

which by inspection holds as long as the mean of the prior is somewhere in the support of the anecdotes.

Thus, the prior being centered anywhere in the support of the anecdotes is sufficient for all of our analysis and results to hold.

M. Positive Weight on the Prior

In our model, we have assumed that a customer uses her prior for a route only if she does not obtain any anecdotes for the route; if she receives one or more anecdotes for the route, her estimate is the sample average of the anecdotes, with no weight given to the prior. We can imagine, however, that a customer might assign positive weight to the prior for a route even if she has anecdotes for the route. In this case, the consequences depend on the amount of noise in the prior, but they are relatively predictable. We suppose for the remainder of this appendix that the prior is centered somewhere in the support of the anecdotes.

First suppose that the noise in the prior (i.e., σ) is small. With small σ, the prior estimate will be in the support of the anecdotes with high probability, and thus by Property 1, for $\alpha \geq \mu_S/\mu_F$, the prior for SF will be smaller than any of the FS anecdotes and vice versa for the FS prior and the SF anecdotes. Hence, assigning positive weight to the prior will have a similar effect to that of increasing the sample size, i.e., it will push the equilibrium closer to herding.

Now consider large σ. In this case, it is much less likely that the prior estimate will fall in the support of the anecdotes. Thus, for the same interval of α, assigning positive weight to the prior now makes it more likely that a customer will incorrectly judge the FS system time to be shorter than the SF system time, pushing the equilibrium *away* from herding.

References

Grimmett, G. S. and Stirzaker, D. (2020). *Probability and Random Processes.* Oxford University Press.

Wang, C.-L. (1989). Simple inequalities and old limits. *The American Mathematical Monthly*, 96(4):354–355.