# Reputational Incentives with Networked Customers

Aidan Smith*

November 8, 2021

**Abstract**

We propose a model of reputational incentives for a firm with privately known type whose customers share reviews via a social network. The firm can choose a non-contractible effort level for each customer, which stochastically improves that customer's review. When customers base their purchase decisions on their friends' reviews, firm incentives for effort are stronger if a customer's review will be read by many friends, and if those friends are not too connected (since their beliefs are easier to influence). From the perspective of ex-ante social welfare, this creates a trade-off between providing incentives and generating learning; more connected networks may allow customers to learn more, but remove firm incentives to exert effort when serving them. When effort is sufficiently productive, ex-ante expected total surplus can be higher when the social network has disjoint components. Our results imply that interventions to increase customer review visibility (such as auto-translate features) may be harmful for social and consumer welfare, if firm effort is sufficiently productive.

Keywords: Learning, Network, Reputation
JEL Classification: D82, D83, D85, L14

*Department of Economics, University of Oxford. Email: aidan.smith@economics.ox.ac.uk. I am grateful to Meg Meyer, Stephen Nei, Alex Teytelboym, and seminar participants at the University of Oxford, Cambridge INET, and the Transatlantic Theory Workshop for helpful comments and suggestions. Latest version available at https://sites.google.com/view/aidansmithsite. All errors are my own.

# 1 Introduction

When customers encounter a new firm for the first time, they may be uncertain about its type. For example, diners may be uncertain about whether the chef at a new restaurant is talented. Early customers will base their purchase decisions on their prior beliefs about the firm's type, but later customers can learn about the firm by observing the purchase decisions and experiences of previous customers.

When earlier customers share reviews with later customers, this can create *reputational* incentives for a firm. In a one-shot setting, if effort is chosen after customers have already decided to purchase, a firm would have no incentive to engage in costly high effort. However, if good reviews today will lead to higher revenue tomorrow, a firm may find it optimal to take costly non-contractible actions when serving customers who visit today. For example, a chef might take more care when preparing a meal, or waiting staff might pay more attention to a table.

In many settings, customers do not observe the full history of reviews, but instead learn via word-of-mouth from their friends on social networks. If firms observe customer networks, they should take into account a customer's popularity when choosing whether to exert high effort for them; if a customer has no friends and only visits a firm once, no future customers will observe their experience, so a firm has no incentives for effort when serving them, while if a customer has many friends, a firm may have strong incentives for effort. Firm incentives when serving a customer will depend not just on how connected that customer is, but on how connected their friends are; if their friends are well connected and have already seen many reviews of the firm, incentives will be weaker on average than when their friends are less well informed and so easier to influence.

This paper asks the following questions: how does customer social network structure determine firm incentives for effort? Which social networks and customer orders should a social planner design if they seek to maximise societal welfare? Our contribution lies in combining a model of reputational incentives with networked learning, and our analysis generates both testable predictions and policy implications.

Our research questions are increasingly important at a time when social networks, platforms, and review aggregators (e.g. Instagram, Etsy, Tripadvisor) all have significant power to influence customer access to past reviews; algorithms can both hide reviews we do want to see (by determining when we see friends' posts) and show us reviews we didn't ask for (via promoted posts). Features such as 'auto-translate' can allow customers to learn from reviews that would otherwise be costly/cumbersome to process. Regulators often lack the knowledge of how platforms manipulate review visibility or the ability to intervene (for example, a regulator cannot observe or control Facebook's algorithms). We provide an insight into how and why a platform or firm might want to manipulate review visibility by exploring conditions under which increasing review visibility unambiguously improves social welfare, and when it does not.

We address our research questions by building a simple model in which a firm with a privately known type serves customers located on a social network, who visit according to some exogenous order. The firm charges a fixed price, and if a customer purchases, the firm chooses whether to exert low or high effort when serving them. The firm is either a commitment 'low-skill' type, who always provides bad service (exerting low effort), or a 'high-skill' type, who sometimes provides good service. Customers are willing to pay more if the firm is high-skill, and are willing to pay more when a high-skill firm exerts high effort. Customers observe the purchase decisions and quality of service provided to their neighbours on the social network, and update their beliefs about the firm's type via Bayes' rule.

For a given network and customer order, we explore what outcomes can be supported in equilibrium. Our main prediction is that persistent good reviews for a firm over time are more likely when customers observe a limited sample of past reviews, rather than the full history, as illustrated in Proposition 3. This arises because when customers observe the full history of reviews, firms can stop exerting high effort following early success and 'coast' on the back of their good reputation. In contrast, when customers networks are less connected, news of early success spreads less widely, meaning firms have less incentive to shirk following early good reviews.

While we find there are equilibria with persistent high effort when customer networks are less connected, we also find for some networks there may be equilibria featuring *herding*, in which the firm exerts high effort for at most the first customer, and all remaining customers purchase if the first customer wrote a good review. In particular, Proposition 1 shows that when customers are located on a directed rooted tree, if customers further from the 'root' move later, there is an equilibrium in which only the 'root' receives high effort, and all remaining customers purchase if the 'root' writes a good review (either because they observe the 'root', or because they copy their friends). Customers who copy their friends' purchase decisions reason as follows: 'if my friend is sufficiently well informed, her decision to visit a restaurant reveals *she* must have learned the firm has had sufficiently many good reviews, so even if she has a bad experience herself I should ignore this and copy her purchase decision'.

Proposition 2 gives conditions under which equilibria are unique, but in general we find multiple equilibria for a given network and customer order. Multiple equilibria arise because equilibria have a feature of 'self-fulfilling prophecy': if customers expect a firm to be exerting high effort, they judge the firm more harshly if it receives bad reviews, which in turn can mean the firm is motivated to high effort. In contrast if they expect the firm to exert low effort, they are less sceptical of its skill level when it receives bad reviews, which can in turn mean the firm chooses not to exert effort. Our results suggest that customers may be better off when they are more optimistic about unknown firms, in the sense that for a given network, if customers expect high effort as often as possible in equilibrium, average customer payoffs may be much higher than if customers expect low effort as often as possible in equilibrium.

We investigate which social networks and customer orders a planner looking to maximise expected social welfare would design. We show that a social planner can strengthen firm incentives by having reviewers visit simultaneously, because when reviewers visit sequentially, a firm may quit exerting high effort early if initial reviews are good enough. We show that if a planner wants to induce high effort for as many customers as possible, providing the strongest possible aggregate incentives to the firm does not in general show all later customers all past reviews; Proposition 5 shows incentives are strongest when later customers see relatively few reviews because the probability that each one of those reviews might determine their decision is larger. Proposition 8 shows that a consequence of this feature of incentives is that when high effort is sufficiently productive, efficient networks will not show customers the full history of reviews, and may have disjoint components. We argue this implies that when reputational incentives are strong, features such as 'auto-translating reviews' may harm rather than benefit total welfare.

The paper is structured as follows; in section 2, we set out the model. In section 3, we discuss what outcomes can be supported in equilibrium for a given social network and fixed order of customers. In section 4, we discuss which networks and customer orders maximise social welfare. Section 5 concludes and discusses how our analysis could be extended to more general settings. In the remainder of this section, we discuss how our paper relates to the literature.

**Related Literature** Our paper's contribution lies in combining a model of reputational incentives with a model of social learning on a network. Our model has some similarities to Mailath and Samuelson (2001)'s model of reputations, although in contrast to their infinite horizon model in which a firm sets a price each period, we consider a firm with finite life span and a fixed price, so in our setting reputational incentives for effort are at the *extensive* margin and will eventually diminish with time. The main distinction of our model is we restrict customers to only observe their neighbours on a social network; in our model, the *identity* of a customer will matter for a firm's effort choice (in other words, we allow *asymmetry* in how much of a game's history customers can observe).

The reputations literature has explored various settings in which customers are restricted to observing only a subset of the history of play. Pei (2021) provides a useful discussion of related literature, and considers a setting with observational learning where customers observe a bounded subset of the history of play between a long-run player and a sequence of short-run players. He shows that allowing a customer to sample a bounded number of previous customer's decisions *in addition* to the history of firm actions can be harmful, in the sense that when customers can also observe previous customers' decision, firms are only guaranteed min-max rather than Stackelberg payoffs in equilibrium. This has a similar flavour to our results on herding; in our model, when customers base their decisions on previous customers' decisions rather than on the outcomes of previous firm effort, this removes firm incentive for effort and can lead to significantly lower equilibrium payoffs.

In contrast to the reputations literature we relax two common assumptions. Firstly, we allow customers to ex-ante differ in information quality for reasons other than the date; using the language of networks allows us accommodate ex-ante *asymmetry* in what customers could observe even if they visited at the same date. Secondly, we do not restrict customers to visiting a firm sequentially, and we show there can be significant welfare gains from allowing customers to visit simultaneously, because this can strengthen incentives for the long-run player.

Both the reputations and social learning literatures commonly fix the *order* in which customers visit[1], and explore different observation rules for customers visiting sequentially. In contrast, we allow for both the customer order *and* observation structure to vary. We note that in the absence of discounting or a strategic long-run player, interventions on customer orders can be mapped to interventions on customer social networks with sequential visits. For example, Sgroi (2002) shows that in a social learning setting where customers observe the full history of play, average payoffs are higher when some later customers are moved to visit in the first period (as when herding begins, herds are better informed). We note that absent discounting the effect Sgroi identifies does not depend inherently on *timing*; Acemoglu et al. (2011) develop a similar insight using the language of networks in a setting with one customer per period, where the first $n$ customers do not observe one another, and the remainder can observe the full history. In our setting, however, interventions on timing are *not* in general equivalent to interventions on networks, because of their effect on firm incentives; a firm forced to choose efforts simultaneously for two customers faces a different problem to a firm who can observe the outcome with the first customer before choosing effort for the second.

Our model of learning links to the social learning literature following Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) (see Golub and Sadler (2016) for a recent survey of models of learning on social networks). Like Acemoglu et al. (2011) and Board and Meyer-ter-Vehn (2021), we focus on Bayesian learning restricted to learning from neighbours on a social network, but unlike in standard social learning models, we do not endow customers with access to a private signal; customers are not born with some inherent private information about firm type, and all 'signals' in our model are the endogenous result of firm effort choice. Board and Meyer-ter-Vehn (ibid.) also feature

[1]A common assumption is that customer visit dates are drawn from some continuous distribution on the unit interval; we note here that absent discounting, this assumption is equivalent to assuming that customers visit the firm sequentially, one per period.

endogenous information, but in their setting, this is the result of customers trialling a product before purchase, while in our setting information is generated by purchases themselves. Like Board and Meyer-ter-Vehn (2021), we focus on learning in the short-run; we are interested in outcomes when a firm serves each member of a finite set of customers exactly once, rather than investigating whether asymptotic learning occurs.

Combining a model of learning on a network with a model of reputational incentives allows us to explore a key trade-off when incentives are reputational and customers observe endogenous signals; observation structures which encourage efficient learning may provide weaker incentives for effort. Our results on efficient networks highlight that policy conclusions drawn from social learning models may be inappropriate when firms can exercise discretionary effort; in a social learning model, it is always optimal ex-ante to allow the final customer to visit a firm to observe the full history; when firms have reputational incentives, this is no longer true (ex-ante, thought it may sometimes be true ex-post), because making the final customer better informed can remove firm incentives for effort for earlier customers. As such, our paper shows the importance of studying reputational incentives and social learning *together*, and suggests significant scope for future work studying more general reputation problems with networked short-run players.

Our questions also have links to the literature on targeting interventions at the most influential agents on networks (for example, targeted information seeding or subsidies). For example, Galeotti and Goyal (2009) study a monopolist who can exert costly effort to influence networked customers via advertising, when customers can learn either directly via advertising or via word-of-mouth communication with neighbours. While we address conceptually similar questions, our model differs in several respects; instead of information spillovers communicating product existence, in our model, they are informative about an underlying type for the firm. We also focus on the *dynamics* of firm effort choice; our firm chooses an individual effort level for each customer, and can condition its effort choices on the history of reviews in previous periods.

# 2 Model

A long-lived firm $F$ serves $N$ customers $\mathcal{C} = \{C_1, .., C_N\}$ across $T \geq N$ periods. The firm has a privately known type $\theta \in \{0, 1\}$, and customers initially have a prior belief that $P(\theta = 1) = \pi$. If $\theta = 0$, the firm is a 'low-skill' commitment type who will always provide customers with bad service; if $\theta = 1$ the firm is a 'high-skill' type who is capable of providing good service.

Each customer visits the firm once, takes a purchase decision, and receives a gross payoff equal to the quality of service received. Each customer is allocated a period in which to visit accordingly to a publicly observed order $\succeq_\alpha : \mathcal{C} \to \{1, ..., T\}$. We will write $C_1^t$ to denote that $\succeq_\alpha$ allocates customer $C_1$ to visit in period $t$. Customers are located on a commonly known network $\mathcal{N}$, and observe the purchase decisions and experiences of their earlier neighbours on $\mathcal{N}$; we write $N(C_k)$ to denote the set of $C_k$'s neighbours on $\mathcal{N}$.

In period $t$, the firm serves all $\mathcal{C}^t = \{C_i^\tau \in \mathcal{C} | \tau = t\}$. For each $C_i^t$, $F$ and $C_i^t$ play the following stage game:

1. $C_i^t$ chooses to purchase ($a_i = 1$) or not purchase ($a_i = 0$) at fixed price $w$ (observed by $F$).

2. $F$ privately chooses effort level $e_i \in \{0, 1\}$.

3. Quality of service $s_i \in \{u_0, 0, 1\}$ is realised. $C_i^t$ receives payoff $s_i - wa_i$, $F$ receives payoff $a_i w - ce_i$, where $c > 0$.

4. $a_i$ and $s_i$ are observed by all $C_j \in N(C_i^t)$, and by $F$.

The firm serves simultaneous customers separately, so can choose different efforts for $C_i^\tau$ and $C_j^\tau$, even though both customers are served simultaneously in period $\tau$. The firm chooses each effort $e_i$ to maximise their discounted sum of expected future payoffs, with discount factor $\delta \in [0, 1]$.

**Firm Technology** The quality of service customers received depends upon their purchase decision, firm effort choice and firm type. If customers do not purchase ($a_i = 0$), their quality of service is just their outside option payoff, which is $s_i = u_0$. If customers do purchase ($a_i = 1$), they either receive good ($s_i = 1$) or bad ($s_i = 0$) service[2] with probabilities $P(s_i = 1|e_i, \theta) = p_{e_i\theta}$ and $P(s_i = 0|e_i, \theta) = 1 - p_{e_i\theta}$.

We make the following assumptions on firm technology: if $\theta = 0$, $F$ is a 'low-skill' type who always provides bad service: $P(s_i = 0|\theta = 0) = 1$. Since effort is costly, it follows that if $\theta = 0$ the firm will always exert low effort. If $\theta = 1$, $F$ is a 'high-skill' type with $p_{01} < p_{11} \leq 1$. As the firm's effort choice is only relevant if $\theta = 1$, going forward we will take $e_i$ to mean the effort choice of $F$ when serving $C_i$ given $\theta = 1$.

We assume $0 < p_{01}$: good service is more likely from a high-skill form than low-skill firm regardless of firm effort choice. In other words, customers prefer a good chef putting in low effort to a bad chef. We make this assumption to guarantee that a reputation is valuable. If customers were indifferent between being served by a high-skill, low-effort firm and a low-skill firm, with a finitely lived firm we would be able to unravel any equilibria featuring high effort; final period customers would know they will receive low

[2] Given how we define quality of service, $s_i$ is a sufficient statistic for $a_i$; as such, the role of allowing customers to observe one another's purchase decisions is purely expositional. We also note that if customer order is commonly known, we do not actually need customers to *observe* non-purchases; in other words, observing $s_i | s_i \in \{0, 1\}$ is a sufficient statistic for $s_i$ and $a_i$; non-purchases can be inferred from a lack of reviews.

effort, and given this, would not care about the firm's type, meaning the firm would have no reputational incentives in the penultimate period, etc[3].

Under our assumptions, only a high-skill firm is capable of generating good service. This makes the customer inference problem simpler; observing one neighbour receive good service is conclusive evidence that the firm is high-skill (while observing only bad service is inconclusive since the firm could either be low-skill or an unlucky high-skill firm). We make this assumption for tractability, but our insights generalise easily to settings in which low-skill firms have a low probability of providing good service. In particular, results will be robust in the following sense; for almost all parameterisations, there exists some $\bar{\epsilon} > 0$ such that if a strategy profile can be supported in PBE for a given network-order pair, when $P(s_i = 1|\theta = 0)$, it can also be supported in PBE when $P(s_i = 1|\theta = 0) = \epsilon$ for $\epsilon < \bar{\epsilon}$. Roughly, because our setting is discrete, for almost all parameterisations it is not crucial that good news is conclusive; it is sufficient that it is merely overwhelming evidence in favour of a high type. However, the model is much less tractable with general networks when $P(s_i = 1||\theta = 0)$ is bounded away from 0.

**Equilibrium Concept** Our solution concept is Perfect Bayesian Equilibrium (PBE). This requires that for any history of play, firm effort choices are a best response to customer strategies, customer purchase decisions are a best response to firm strategies given customer beliefs about firm type, and wherever possible customers derive their beliefs about firm type use Bayes's rule. As our setting is finite, the concept of PBE is well defined and we can apply standard textbook definitions. Below, we offer a formal definition of PBE our setting.

Define the history of the game at the beginning of period $\tau$ as $H_\tau = \{(a_i, s_i)|C_i \in \mathcal{C}^t\}_{t=1,\dots\tau-1}$, and define $\mathcal{H}_\tau$ as the set of all possible $H_\tau$s. An information set $I_i^\tau$ for customer $C_i$ in period $\tau$ is $I_i^\tau = \{(a_j, s_j)|(a_j, s_j) \in H_\tau, C_j \in N(C_i)\}$, for some history $H_\tau$; define $\mathcal{I}_i^t$ as the set of all possible $I_i^\tau$s. A behavioural strategy for each customer $\rho_i : \mathcal{I}_i^t \to [0, 1]$ assigns a probability of purchasing (setting $a_i = 1$) to each possible information set for customer $C_i$ at time $t$. A behavioural strategy for the firm $\sigma : H_t \times \mathcal{C}^t \to [0, 1]^{|\mathcal{C}^t|}$ assigns an effort choice[4] to each customer who visits in period $t$, for each history of the game at the start of period $t$. Beliefs for customer $C_i$ in period $t$ at information set $I_i^t \in \mathcal{I}_i^t$ are given by $\mu_i^t : \mathcal{I}_i^t \to [0, 1]$, where $\mu_i^t(I_i^t) = P(\theta = 1|I_i^t, \bar{\sigma}_i)$, $\bar{\sigma}_i$ is $C_i$'s conjecture about

[3]The important point here is not that the firm is finitely lived; the same argument can hold with an infinitely lived firm. The crucial point is that for reputational incentives to exist when incentives are for a strategic type to *distinguish* themselves from a bad type, rather than to *imitate* a good type, it must be that once the strategic type has successfully distinguished themselves, customers still care about their type.

[4]Without loss of generality, we assume a firm does not condition future effort choices on past effort choices. As customers do not observe effort choices, it is sufficient for the firm to condition future effort choices on past quality of service; if the firm shirks in period $t = 1$ and gives bad service to all customers, their problem in period $t = 2$ is identical to what their problem would have been if they had exerted high effort in period $t = 1$ but been unlucky and still achieved bad service.

the firm's strategy $\sigma$, and $\mu_i^t$ is computed via Bayes' rule at information set $I_i^t$ wherever possible given conjecture $\bar{\sigma}_i$.

Formally, a Perfect Bayesian Equilibrium (henceforth equilibrium) is a strategy for each customer $\rho_i$ $\forall i = 1, ..., N$ and a strategy for the firm $\sigma$, such that $\rho_i$ maximises $C_i$'s expected payoff at each information set $I_i^t$ given conjecture $\bar{\sigma}$, and $\sigma$ maximises the firm's expected discounted payoff at each history $H_t$ given $\rho_i$s, and $\forall i = 1, ..., N$, $\bar{\sigma}_i = \sigma$. We will focus on pure strategy equilibria.

We note that given our firm technology, with short-lived customers, off-path beliefs will not play a role in supporting outcomes in equilibria where all customers have a positive probability of purchasing from the firm (though they may play a role in supporting equilibria where some customers do not purchase). We are largely interested in equilibria with positive probabilities of trade between all customers and the firm; as such, we will typically omit discussion of off-path beliefs for the sake of brevity.

## 2.1 Modelling Assumptions

In the remainder of this section, we discuss some of our modelling assumptions.

Throughout, we interpret the network $\mathcal{N}$ as a social network, but its sole role in our analysis is to determine which customers observe one another. As such, the reader is free to treat the language of networks as purely a modelling tool capturing an underlying observation structure (for example, an organisational hierarchy, or a rule determining from how many previous periods a customer can observe reviews). Since we restrict attention to fixed endogenous orders, our analysis is also unchanged if the network is directed or undirected, in the sense that for $C_i^t$ and $C_j^{t+s}$, because $C_i^t$ moves first, it is unimportant whether $C_i^t$ observes $C_j^{t+s}$, since $C_i^t$ has no more actions left to take after period $t$.

While we assume that each customer visits the firm only once, it is possible to interpret distinct customers in our model as multiple instances of the *same* customer acting myopically. For example, if $N(C_i^t) \cup C_i^t = N(C_j^{t+k}) \cup C_j^{t+k}$, then $C_i^t$ and $C_j^{t+k}$ share the same friends and we can interpret $C_i^t$ and $C_j^{t+k}$ as the same customer who visits both in period $t$ and in period $t+k$. In particular, if $\mathcal{N}$ is a complete social network and $\succeq_\alpha$ is a total order on $\mathcal{C}$ (in other words, one customer visits in each period), interpreting the model as one myopic customer visiting the firm $N$ times is equivalent to interpreting the model as $N$ customers each visiting the firm once. Note the model *would* differ with repeat customers who were forward looking; relative to the myopic case, customers might choose to purchase from a firm not because they expected good service, but in order to gather information to inform future purchase decisions[5].

[5]This incentive for customers to gather information would be unlikely to be a significant concern if the number of total customers is large relative to the number of visits each customer makes. For example, with e.g. 50 customers each visiting twice, information gathering incentives on first visit are likely to be minimal if the social network is sufficiently connected.

We consider a firm with two possible types; a strategic type which faces a meaningful effort choice problem, and a simple type who always provides bad service. Incentives arise because the strategic type would like to distinguish themselves from the simple type in the eyes of customers. More generally, we can separate reputational incentives in games into two types of incentive: incentives to distinguish oneself from a bad type, and incentives to imitate a good type (we can think of these incentives as analogues of incentives in separating and pooling equilibria in signalling games).

Many properties of incentives are general to both categories, but there are also important distinctions. For example, in infinite time horizon settings where customers observe the full history, incentives to distinguish oneself are inherently less persistent than incentives to imitate. Incentives to distinguish are not persistent because as time passes, firm type will become statistically identified on the equilibrium path for almost all histories, rendering firm effort choice irrelevant. In contrast, incentives to imitate can be very persistent, because if in equilibrium a strategic type takes exactly the same actions as some commitment (Stackelberg) type, customers do not learn more about the firm's type as time passes, so incentives are time invariant. Which form of incentives is more appropriate to focus on in modelling will depend upon the setting, but we note that empirically, businesses *do* differ in inherent quality level, many new businesses *do* fail when they acquire bad reputations, and many new businesses *do* succeed in building good reputations.

Throughout, we take the order of customers visits as exogenous from the firm's perspective. In some settings, this will be a reasonable assumption; for example, the order in which students take driving lessons is likely to be determined by the order of their birthdays. More generally, understanding equilibrium outcomes for a fixed order of customers is an important first step to understanding equilibria in settings in which customers or firms can influence the order of visits; for example, settings in which customers need to book restaurant tables in advance, or in which restaurants can specifically invite certain customers to their opening.

We are modelling the price of the firm as fixed for the whole game. An alternative approach to reputational incentives is to allow the firm to choose a price for each customer, or each period. Under this alternate specification, reputational incentives would be conveyed via a different channel. If pricing is fixed, incentives for effort are transmitted via the *extensive* margin; high effort may increase the probability that future customers are willing to buy from the firm at price $w$. If pricing is discretionary, incentives for effort are entirely at the *intensive* margin; the firm charges each customer their ex-ante willingness to pay, extracting all surplus, and all customers participate in equilibrium. Modelling prices as fixed is significantly more tractable and allows customers to draw inference from the purchase decisions of their neighbours. In many settings fixed pricing is also the more realistic assumption in the short run; a new restaurant who wants to build a reputation is unlikely to make minute adjustments to their advertised menu after serving each customer.

# 3 Strategies and Equilibria

In this section, we discuss how customer decisions and firm incentives depend on network structure. We first introduce customers' purchase problems and the firm's effort choice problem, before illustrating how network structure affects customer learning and firm incentives via three examples. We then state results for more general network settings.

## 3.1 Customer Strategy

Each customer will purchase at an information set if their willingness to pay at that information set minus price $w$ is greater than their outside option $u_0$. Customer $C_i$, expecting $e_i = b$, purchases at information set $I_t$ if:

$$u_0 \leq P(\theta = 1 | I_t, \bar{\sigma}_t) p_{b1} - w$$

Customers can learn both from the experiences of their neighbours and from the purchase decisions of their neighbours. If $C_t$ believes that all of their earlier friends purchase with certainty, then $C_t$ will only draw inference from the quality of service received by their neighbours, since the purchase decisions of earlier neighbours are uninformative about the neighbours' information sets. If $C_t$ believes $P(a_{t-k} = 1) < 1$ for some $C_{t-k} \in N(C_t)$, $C_t$ can also draw inference from the purchase decision of neighbour $C_{t-k}$; they can infer if $a_{t-k} = 1$ then $C_{t-k}$ must have seen good news about the firm, while if $a_{t-k} = 0$, $C_{t-k}$ can only have seen bad news. In our setting, because good news is conclusive, if a customer only sometimes purchases, when they do purchase they must be certain that the firm is a high type, since they can only have seen either conclusive good news or inconclusive bad news.

For equilibria featuring trade to exist at all, there are two important information sets at which a customer should be willing to purchase: a customer who has learned the firm to be a high-skill type must be willing to purchase expecting low effort (else we would be able to unravel any candidate equilibrium featuring trade, since the final customer, knowing they would receive low effort, would never purchase, and hence the penultimate customer would never purchase, etc), and a customer who expects high effort must be willing to purchase from the firm under their prior. For incentives for effort to ever exist, it must also be the case that customers are not willing to purchase from a firm they know to be a low type for sure (else reputations do not matter). As such, for the remainder of the paper, we make the following assumption:

**Assumption 1.** *$w \in (-u_0, p_{01} - u_0]$ and $w \leq \pi pi_{11} - u_0$.*

We will define a *critical customer* as a customer in an equilibrium who purchases from the firm at some (but not all) of their information sets on the equilibrium path.

**Definition 1.** *A critical customer for a given strategy profile* $(\boldsymbol{\rho}, \sigma)$ *is a customer* $C_i$ *such that ex-ante,* $0 < P(a_i = 1|\boldsymbol{\rho}, \sigma) < 1$.

## 3.2 Firm Strategy

The firm will exert high effort when serving customer $C_i$ if doing so increases their discounted expected revenue by more than the cost of effort, $c$. Formally, when serving $C_i^t$, the firm's incentive compatibility (IC) constraint for high effort is given by:

$$c \leq \sum_{k=1}^{T-t} \delta^k w[P(a_j = 1|e_i = 1) - P(a_j = 1|e_i = 0)]\mathbf{1}_{C_j^{t+k} \in N(C_t)}$$

We can see a necessary condition for exerting high effort is that doing so increases the probability that future customers purchase. If the firm believes that all of $C_i$'s neighbours will always purchase regardless, the firm has no incentive for effort when serving $C_i$ and should set $e_i = 0$. In other words the firm will consider exerting high effort for $C_i$ if and only if $C_i$ is observed by a critical customer. In particular, any customers who are observed by no-one (for example, the final customer) will receive low effort in any equilibrium.

## 3.3 Examples

We now give three examples of networks and customer orders to illustrate the firm and customer problems in our model:

**Example 1: Customers observe their predecessor** In example 1, we have 4 customers who visit sequentially, each of whom observe their immediate predecessor (for example, we could think of a setting where all customers write reviews, and review websites only display a firm's most recent review).

We will focus on the case where if the final customer observes a bad review, they are unwilling to purchase (if all customers who have seen a bad review still purchase, in the firm optimal equilibrium the firm would never exert effort and all customers would purchase regardless). Formally, assume $\frac{\pi(1-p_{01})p_{01}}{\pi(1-p_{01})+(1-\pi)} < w + u_0$.

To understand how incentives and learning operate in this setting, let's suppose that the firm exerts high effort for the first customer: $e_1 = 1$ (we know by assumption that if the firm plans high effort for $C_1$, $C_1$ is willing to purchase), and ask what this implies about equilibrium outcomes.

The first thing we can note is that if the firm exerts high effort for $C_1$, it must be that $C_2$ only purchases if they see a good review from $C_1$; if $C_2$'s purchase decision does not

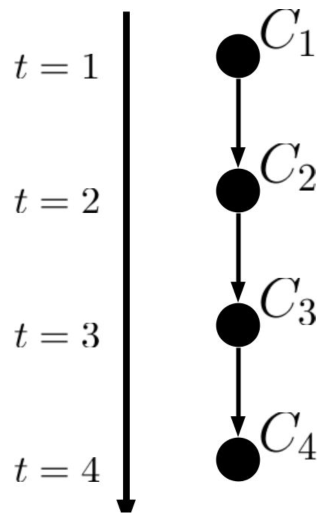


Figure 1: Example 1: Customers observe their predecessor

depend on $C_1$'s review, the firm would be better off setting $e_1 = 0$ as effort is chosen after $C_1$ chooses whether or not to purchase.

The second point to note is that if $C_2$ only purchases when they see a good review ($s_1 = 1$), then $C_3$ can learn from $C_2$'s purchase decision. If they see $C_2$ purchase they learn the firm is a high type for sure, whereas if they see $C_2$ does not purchase, they should revise their beliefs downwards. In other words, although $C_3$ does not observe $C_1$ directly, if $C_2$'s purchase decision depends on $s_1$, then in equilibrium $C_3$ learns the value of $s_1$ anyway.

An important implication of $C_3$ learning via $C_2$'s purchase decision is that the firm should not exert effort for $C_2$. To see this, note that the firm should only exert effort for $C_2$ if $C_2$'s review influences $C_3$'s decision. But we can note that $C_2$'s review can never influence $C_3$'s decision; $C_2$ only purchases if they see a good review, which implies the firm is a high type for sure. As a result, any time $C_3$ sees $C_2$ write a review, they can deduce that the firm is a high type because $C_2$ purchased, and so $C_3$ learns nothing extra from the review's content[6]. In other words, high effort for $C_1$ implies low effort for $C_2$.

Combining these insights, we can reason as follows; $C_3$ learns everything $C_2$ knows in equilibrium, so if $C_3$ also expects low effort from the firm, $C_3$'s problem is identical to $C_2$'s. As a result, if they expect low effort, they should always copy $C_2$'s purchase decision (even if they expect high effort, they should copy $C_2$ if $\frac{\pi(1-p_{11})p_{11}}{\pi(1-p_{11})+(1-\pi)} < w$). Further, given $C_3$ copies $C_2$, $C_4$ should copy $C_3$ (note $C_4$ *knows* they will receive low effort, because they are observed by no-one).

What this tells us is that in example 1, we can find an equilibrium where if $C_1$ receives high effort, all remaining customers purchase only if $C_1$ writes a good review, and the firm

[6]This would not be true if good reviews did not fully reveal the firm to be a high type, but an analogous point holds in more general settings; if the amount $C_3$ can learn from $C_2$'s review is small compared to what they learn from $C_2$'s purchase decision, they should copy the purchase decision regardless of the review contents.

exerts high effort for $C_1$ only. Formally, this equilibrium exists if $c < (p_{11} - p_{01})w \sum_{i=1}^{3} \delta^i$ and $\frac{\pi(1-p_{11})p_{01}}{\pi p_{11}+(1-\pi)} < w + u_0 < \pi p_{11}$, and is unique[7] if $\frac{\pi(1-p_{11})p_{11}}{\pi(1-p_{11})+(1-\pi)} < w + u_0$.

This equilibrium is quite striking. When customers observe only their immediate predecessor, there is an equilibrium in which the firm exerts high effort at most once and all remaining customers base their purchase decisions solely on the review of the first customer. We can see this equilibrium is bad for customers in two ways ways; firstly, each customer only bases their decision on one review (so with probability $\pi(1-p_{11})$ the firm is high-skill but customers fail to buy), and secondly, only one customer receives high effort when the firm is high-skill.

Example 1 illustrates an important point in general settings; herding is possible in equilibrium and is harmful for firm incentives, because if all remaining customers plan to copy their predecessor's purchase decision, the firm can no longer influence purchase decisions, and so should exert low effort for all remaining customers. We discuss how the possibility of herding generalises in Proposition 1.

**Example 2: Two reviewers, one reader** In example 2, we have two 'reviewers', each of whom observes no-one, and one 'reader', who observes both reviewers and visits the firm after them. Customer $C_t$ visits in period $C_\tau$.

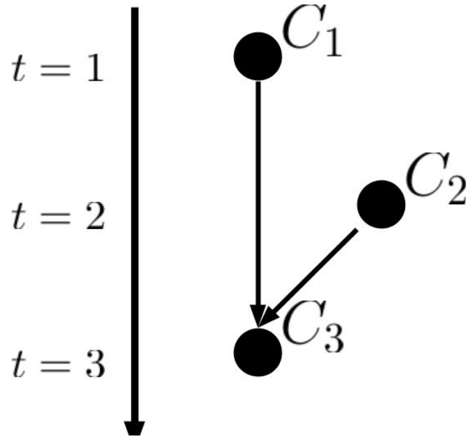


Figure 2: Example 2: Two reviewers, one reader

Again, we will focus on the interesting case where the 'reader' bases their decision on the contents of the reviews. Formally, assume $\frac{\pi(1-p_{01})^2 p_{01}}{\pi(1-p_{01})^2+(1-\pi)} < w + u_0$, to rule out an equilibrium in which the firm exerts no effort for any customer and all customers always purchase. Figure 3 illustrates how equilibrium possibilities depend upon $c$ and $\pi$.

Because good reviews are fully revealing, if $C_3$ bases their decision on the contents of reviews, they should only purchase if they see at least one good review. Stepping back to

[7] If this is not the case, then if $c < (p_{11} - p_{01})w\delta$ there is an additional equilibrium, in which $C_2$ and $C_4$ receive low effort and purchase only if they see a good review, while $C_1$ and $C_3$ receive high effort, and always purchase at any history.

the firm's effort choice for $C_2$, this tells us that the firm should only consider high effort for $C_2$ if $s_1 = 0$; if the firm gets a good review from $C_1$, $C_3$ will learn that the firm is a high type for sure, and so the firm has no more reputational incentives[8] for effort. If $s_1 = 0$, the firm should exert high effort for $C_2$ if $(p_{11} - p_{01})w\delta > c$.

Suppose the firm never exerts high effort for $C_2$. We can argue that in that case, they should also never exert high effort for $C_1$. On an intuitive level, when the firm serves $C_2$ after getting a bad review from $C_1$, the firm understands that it is their last chance to convince $C_3$ to purchase. In contrast, when serving $C_1$, the firm knows they will have a second chance; if they get a bad review from $C_1$, that may not matter if they get a good review from $C_2$. In region $C$ on figure 3, the firm sets $e_1 = 0$, $e_2 = 0$, and $C_1$ and $C_2$ always purchase. In region $E$, the firm would set $e_1 = 0$, $e_2 = 0$ if $C_1$ or $C_2$ purchased, and understanding this, $C_1$ and $C_2$ are too sceptical of firm type so do not purchase.

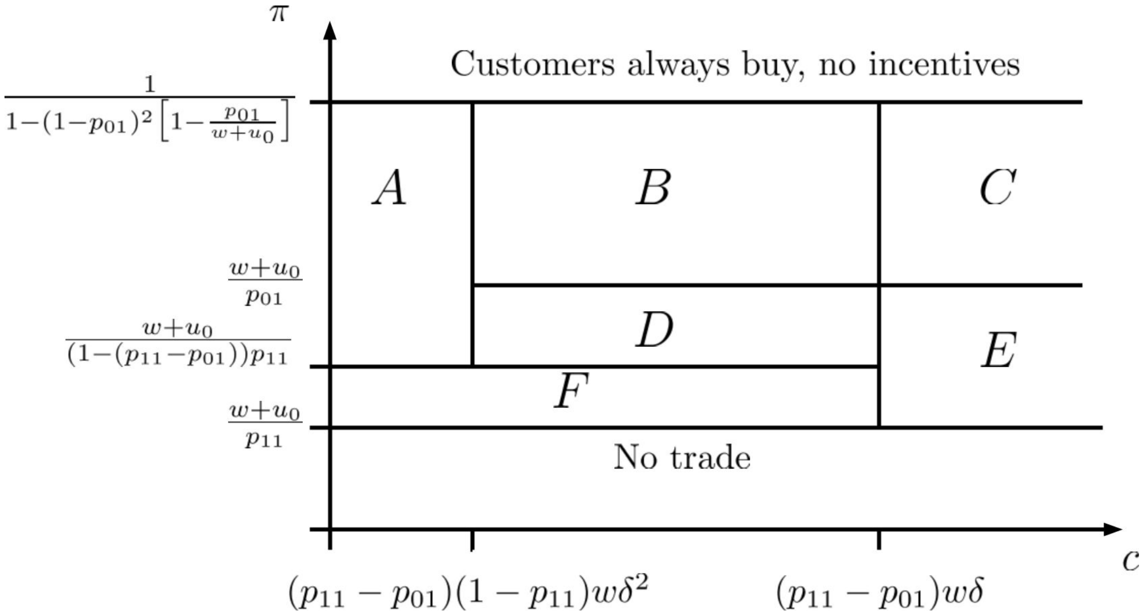


Figure 3: Example 2: Two reviewers, one reader

If the firm is willing to exert high effort for $C_2$, we can ask whether they should exert high effort for $C_1$? Their IC constraint for $C_1$ is now $(p_{11} - p_{01})(1 - p_{11})w\delta^2 > c$. We can see that (even for $\delta = 1$), this is harder to satisfy than their IC constraint for $C_2$ conditional on $s_1 = 0$, because the firm has an incentive to 'wait-and-see' when serving $C_1$, knowing they will get a second chance at a good review when serving $C_2$.

In region $A$ effort is sufficiently cheap that the firm finds the following strategy optimal; exert high effort for $C_1$, and exert high effort for $C_2$ iff $s_1 = 0$, and $C_1$ and $C_2$ have high

[8] With fully revealing good news, if all remaining customers have seen at least one good review, the firm no longer has incentives for effort, but the benefits to the firm of quitting high effort after early success arise more generally. If reviews were only boundedly informative about firm type, then the firm only has reputational incentives in histories where it is still possible to influence a customer's decision; if a customer needs to see four good reviews out of five and the firm has already received four good reviews, the firm no longer has reputational incentives related to that customer.

enough priors to always purchase from the firm.

In contrast in region $F$, $C_2$'s prior is too low to be willing to purchase if the firm plans high effort for $C_1$ (since $C_2$ knows they only get high effort if the firm fails to receive a good review from $C_1$), and so in equilibrium only one customer can purchase from the firm. Effort remains cheap enough that the firm is willing to exert effort if only one customer visits, so in region $F$ in equilibrium one of $C_1$ and $C_2$ always purchases and receives high effort and the other never purchases.

In region $B$ the firm exerts high effort for $C_2$ iff $s_1 = 0$, and does not find it worthwhile to exert high effort for $C_1$ given they know they will have a second chance when serving $C_2$. $C_1$ and $C_2$ always purchase, and $C_3$ purchases if she sees a good review. In region $D$ the unique equilibrium has $C_2$ always purchase and receive high effort, while $C_1$ does not purchase. $C_1$ understands that if they visit knowing $C_2$ also visits in equilibrium they will receive low effort, and their prior is low enough that this deters purchasing, while $C_2$'s prior is high enough that they are willing to visit regardless of whether $C_1$ visits or not.

If we consider expected payoffs, we can see that there are parameterisations where $C_2$'s expected payoff in equilibrium is smaller when effort is cheaper. For example, in region $A$, with cheap effort $C_1$ and $C_2$ purchase and $C_2$ receive high effort with probability $\pi(1 - p_{11})$, while in region $B$ with an intermediate effort cost again $C_1$ and $C_2$ purchase, but now $C_1$ receives low effort and $C_2$ receives high effort with probability $\pi(1 - p_{01})$. In other words, weakening firm incentives for $C_1$ by increasing $c$ can be good for $C_2$ since from the firm's perspective, efforts for $C_1$ and $C_2$ are substitutes (similarly $C_2$ is better off with parameterisations in $D$ than in $B$, so can also benefit from lower $\pi$ for some parameterisations).

If we consider relative payoffs, $C_2$'s expected payoff is higher than $C_1$'s for intermediate costs (regions $B$ and $D$) but lower than $C_1$'s for low costs (region $A$). If effort is very cheap, $C_1$ receives high effort whenever the firm is high-skill, whereas $C_2$ only receives high effort from a high-skill firm with probability $(1 - p_{11})$. If effort is more expensive however, the incentives to 'wait-and-see' become too strong, and $C_2$ is better off than $C_1$; $C_1$ never receives high effort from a high-skill firm, while $C_2$ continues to receive high effort from a high-skill firm with probability $(1 - p_{11})$.

In fact, if effort is cheap $(c < p_{11} - p_{01})(1 - p_{11})w\delta^2)$, $C_2$ would be better off if they could move their visit to the firm to $t = 1$. As they do not observe anyone they derive no informational benefit from visiting the firm in $t = 2$, and when effort is cheap they are less likely to receive high effort than a $t = 1$ customer because the firm can condition $t = 2$ effort choices on $t = 1$ performance. Visiting in $t = 1$ removes the firm's ability to 'wait-and-see', and so for cheap effort, high-skill firms would exert high effort for both $C_1$ and $C_2$ if both visited at $t = 1$. This illustrates an important more general point; even for a fixed network, the timing of customer visits is important for expected payoffs; it is good to be an early reviewer when effort is cheap, and a later reviewer when effort is expensive.

Our second example demonstrated how firm incentives for effort can vary according to the timing of reviewer visits. Our third example illustrates how firm incentives for effort depend on how many customers observe a review, and how well informed those customers are.

**Example 3: Shared and Exclusive Readers** In Example 3, we have 3 'reviewers' who visit at $t = 1$, and 3 'readers' who visit at $t = 2$. $C_1$ and $C_2$ are both observed by $C_4$ and $C_5$, and $C_3$ is observed by $C_6$.

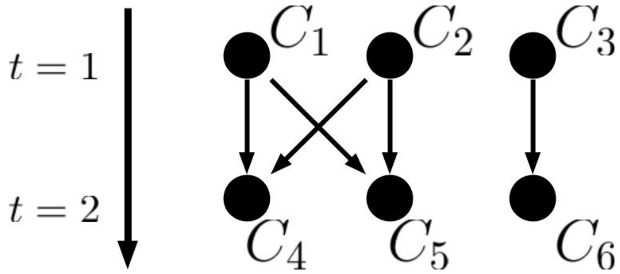


Figure 4: Example 3: Shared and Exclusive Readers

For all readers to base their decisions on the contents of reviews in equilibrium, we need $\frac{\pi(1-p_{01})p_{01}}{\pi(1-p_{01})+(1-\pi)} < w + u_0$, to rule out an equilibrium in which the firm exerts no effort for any customer and all customers always purchase. Note however that if $\frac{\pi(1-p_{01})p_{01}}{\pi(1-p_{01})+(1-\pi)} < w + u_0 < \frac{\pi(1-p_{01})^2 p_{01}}{\pi(1-p_{01})^2+(1-\pi)}$, it could be that in equilibrium $C_6$ always purchases and $C_3$ receives low effort, while $C_4$ and $C_5$ base their decisions on review contents. In other words, if customers have relatively high priors, one bad review may not deter them from purchasing; customers may need to see many bad reviews to be deterred from purchasing from the firm in equilibrium.

We are interested in comparing customer learning from low-effort reviews to learning from high-effort reviews, so we will also assume for this example that customers are willing to purchase from the firm under their prior regardless of effort level ($w + u_0 < \pi p_{01}$).

If $\frac{\pi(1-p_{01})p_{01}}{\pi(1-p_{01})+(1-\pi)} < w + u_0$, the firm knows that 'readers' will only purchase if they have seen at least one good review. The firm's effort choice problem for $C_6$ is simple: if $(p_{11} - p_{01})w\delta \geq c$, exert high effort, otherwise exert low effort.

Turning to $C_1$ and $C_2$, we can see that the firm's effort for one reviewer will affect their incentives for effort for the other. Note that this is not moral-hazard-in-teams; as the firm chooses $e_1$ and $e_2$ simultaneously to maximise a single objective, it can choose the co-operative solution and there is no hold-out problem. Suppose the firm plans high effort for $C_1$. Then the firm should exert high effort for $C_2$ if $(p_{11} - p_{01})(1 - p_{11})2w\delta \geq c$, and should set $e_2 = 0$ otherwise. The firm's effort choice problems for $C_1$ and $C_2$ are symmetric, so if $(p_{11} - p_{01})(1 - p_{11})2w\delta \geq c$, the firm should exert high effort for both.

If $(p_{11} - p_{01})(1 - p_{01})2w\delta \geq c > (p_{11} - p_{01})(1 - p_{11})2w\delta$, the firm only finds it optimal to exert high effort for one of $C_1$ and $C_2$. If $(p_{11} - p_{01})(1 - p_{01})2w\delta < c$, the firm finds high effort for neither customer optimal.

Comparing incentives for effort for $C_1$, $C_2$, and $C_3$, we can see that if $p_{11} < \frac{1}{2}$, then if $C_3$ receives high effort in equilibrium, $C_1$ and $C_2$ must also receive high effort in equilibrium. In contrast if $p_{11} \geq \frac{1}{2}$, it is possible in equilibrium that $C_3$ receives high effort, but either $C_1$ or $C_2$ does not. This illustrates an important point; the strength of firm incentives for a reviewer depends not only on how many later customers will read their review, but also on how many other reviews those later customers read. If a single high effort review is very informative ($p_{11} > \frac{1}{2}$), a firm may not find it worthwhile to exert high effort for multiple reviewers, since it is unlikely a second high effort review will change the mind of a customer who has already seen one high effort review. The intuition is particularly clear in the extreme case as $p_{11} \to 1$; if high-effort guarantees a good review, then the firm should never exert high effort for both $C_1$ and $C_2$, since exerting high-effort once is sufficient to guarantee $C_4$ and $C_5$ observe a good review.

From the perspective of customer expected payoffs, we can say the following: if effort is very cheap then all $t = 1$ customers have the same expected payoff, and $C_4$ and $C_5$ have higher expected payoffs than $C_6$ as they observe more reviews. If the cost of effort is very large, then again all $t = 1$ customers have the same expected payoff as none receive high effort, and $C_4$ and $C_5$ have higher expected payoffs than $C_6$ as they observe more reviews. For intermediate costs of effort, however, if $p_{11}$ is sufficiently large (high effort reviews are sufficiently informative), it may be that $C_3$ has a greater expected payoff than $C_1$ and $C_2$ and $C_6$ has a greater expected payoff than $C_4$ and $C_5$, if the firm chooses high effort for $C_3$ but not for $C_1$ or $C_2$. Note in particular that if one high effort review is more informative than two low effort reviews, for intermediate costs of effort and $p_{11}$ sufficiently large, $C_4$ would be better off ex-ante if they could commit to not reading $C_2$'s review, since this would strengthen firm incentives when serving $C_1$ and improve $C_4$'s information quality.

Example 3 illustrates an important general point; which customers are best off for a given order and given network depends in a non-trivial way on the parameters of the firm's moral hazard problem. When the firm finds effort cheap and reviews are relatively uninformative, in general customers benefit from being more connected. In contrast, if the firm finds effort more expensive and reviews are relatively informative, uninformed customers (reviewers) may be better off being observed by a few 'captive' readers than many shared readers, and readers may be better off committing to reading only a few reviews from reviewers who receive high effort. We return to this further in section 4 when we discuss efficient networks.

## 3.4 General Networks

We now turn to state three results for general networks. Firstly, we show that the 'herding' behaviour in Example 1 is possible in much more general network settings. Secondly, we give conditions of equilibria to be unique, but argue in general we should expect multiple equilibria for a given network and customer order. Finally, we compare networks in which customers observe only the previous customer to networks in which customers observe the full history, and show that if we select the equilibria which feature the most high effort, total effort can be higher when customers only observe the previous customer than when customers observe the full history.

**Herding:** Example 1 (where customers observe only their predecessor and visit sequentially) showed that *herding* is possible in equilibrium, with all customers after some date copying the purchase decision of their predecessor regardless of reviews. We can extend this insight to *any* network where customers visit sequentially and all customers are connected by a path on the network to the first customer:

**Proposition 1.** *If one customer $C_t$ visits in each period $t$, $c < (p_{11} - p_{01})w \sum_{i=1}^{N-1} \delta^i$, $\frac{\pi p_{11} p_{01}}{\pi p_{11} + (1-\pi)} < w + u_0 < \pi p_{11}$, and all $C_\tau$ for $\tau \geq 1$ are connected to $C_1$ via a path on $\mathcal{N}$, there exists an equilibrium in which $e_1 = 1$, $C_1$ always purchases, and all remaining customers $C_\tau$ purchase iff $s_1 = 1$, and we have $e_\tau = 0$ for $\tau \geq 0$.*

Note that any network where all customers observe the full history of reviews satisfies this requirement. This requirement is also satisfied by networks which are directed rooted trees (formally, networks satisfying, for $t > 1$, $|\{C_{t-k} | C_{t-k} \in N(C_t), k > 0\}| = 1$).

This result gives a sufficient condition for there to exist an equilibrium in which all customers base their purchase decisions on the review of the first customer. We can note that if $C_2$ bases their purchase decision on $C_1$'s review, and all remaining customers are connected via a directed path to $C_2$, we can also restate the result as giving a sufficient condition for all remaining customers to copy $C_2$'s purchase decision. This gives us the following corollary:

**Corollary 1.** *Suppose one customer $C_t$ visits in each period $t$, and all customers who visit after $\tau$ are connected by a path to $C_\tau$, and are not connected to any $C_j$ for $j < \tau$. Then if $C_\tau$ is a critical customer, there exists an equilibrium in which all remaining customers copy $C_\tau$'s purchase decision.*

Our corollary highlights that even on networks where not all customers are connected via a path to the first customer, we can identify subgraphs of a social network on which herding *is* possible. We can take the following approach to characterising *an* equilibrium for any general network:

1. Fix customers' conjectures $\bar{\boldsymbol{\sigma}}$ about firm strategy.

2. Identify the first critical customer using in-degree. All uncategorised customers before them always purchase.

3. Identify the set of customers who are connected by a path to that critical customer (but not to any earlier/contemporaneous customers). There exists an equilibrium in which these customers copy the critical customer.

4. If not all customers are categorised, identify the earliest uncategorised customer and repeat above steps.

5. Having identified all critical customers, identify the customers who are observed by critical customers. These customers are the only candidates to receive high effort.

For many networks, using the approach described above can significantly simplify the firm's effort choice problem by identifying customers and links which are irrelevant for a firm's effort choice problem.

For example, consider the case of a complete social network with $C_t$ visiting in period $t$, in which given their conjectures the first 10 customers always purchase, and customers $C_{11}$ and later only purchase if they have observed conclusive good news about the firm's type. We can see that given customer conjectures, the firm's effort choice problem is in fact equivalent to the problem of a firm facing the following network; $N(C_1), ..., N(C_{10}), N(C_{12}), ..., N(C_T) = C_{11}$, $N(C_{11}) = \mathcal{C} \setminus C_{11}$. In other words, with this conjecture, the complete social network is equivalent[9] to a star network in which $C_{11}$ is the centre of the star; since all customers after $C_{11}$ only purchase in histories in which $C_{11}$ also purchases, it is equivalent to have them only observe $C_{11}$, and given that customers $C_1, .., C_{10}$ always purchase, it is equivalent to delete all links amongst them and have them observed only by $C_{11}$. Phrased more generally, links among customers who always purchase are payoff irrelevant, and in any setting in which customers herd, it is payoff-equivalent to link all herd members to the 'herd leader' and delete all other links of the non-leaders. We state a stronger result which captures this reasoning in proposition 4.

Our results on the possibility of herding highlights how much the strength of reputational incentives can differ when incentives are at the *extensive* rather than *intensive* margin. In our model, if each customer only observes the previous customer, two consecutive customers cannot both receive high effort, because when purchase decisions are informative about firm type, this removes incentives for consecutive effort. In contrast, if a firm could set prices, they would always charge customers their willingness to pay, and customers would always purchase. As a result, while a fixed price firm cannot exert high effort for two consecutive customers, it is possible that a price *setting* firm could exert high effort for all but the final customer, because for a price setting firm, customer purchase decisions are not informative about firm type.

[9] Note that we only have equivalence *given* customer conjectures; for other conjectures about firm strategy, equilibria may differ significantly between the two networks.

**Uniqueness:** Multiple equilibria arise in general in our setting because customer willingness to pay depends upon whether customers believe a firm has been exerting high effort (and whether they expect the firm to continue to exert high effort). Customers judge a firm which receives bad reviews more harshly if they believe it has been trying, but are willing to pay more if they think they will receive high effort if the firm is high skill. For equilibria to be unique, we must remove the possibility of self-fulfilling prophecies, where the firm's equilibrium effort choices depend upon customers' conjectures about firm effort.

We can define the following two thresholds for price $w$:

$$\bar{t} = \max\{t | t \in \mathcal{N}, \frac{(1-p_{01})^t \pi}{(1-p_{01})^t \pi + (1-\pi)} p_{11} \geq w + u_0\}$$
$$\underline{t} = \max\{t | t \in \mathcal{N}, \frac{(1-p_{11})^t \pi}{(1-p_{11})^t \pi + (1-\pi)} p_{01} \geq w + u_0\}$$

$\underline{t}$ is the largest in-degree a customer can have for which we can guarantee that they always purchase in any equilibrium; $\bar{t}$ is the largest in-degree a customer can have such that it is possible they always purchase in *some* equilibrium[10]. We can use these thresholds to give the following result:

**Proposition 2.** *If $\bar{t} = \underline{t}$, and each customer $C_t$ visits in period $t$, actions on the equilibrium path in a pure strategy PBE are unique for almost all parameterisations.*

If $\bar{t} = \underline{t}$ the initial set of customers who purchase only if they have observed a good signal does not depend on customer conjectures about firm effort choice. Hence we can eliminate self-fulfilling prophecies in equilibria; since customers strategies are independent from conjectures about firm type, firm strategies should also be independent from customer conjectures. Further, if we only have one customer per period, given our discrete setting, firm optimal strategies will generically be unique for fixed customer strategies (we can reason that with all effort choices taken sequentially, for almost all parameterisations the firm will have a unique optimal action when serving the final customer at each history, and when serving the penultimate customer, etc.).

Note we cannot guarantee equilibrium uniqueness in a general setting when customers visit the firm simultaneously, because for many parameterisations a firm may be indifferent between several strategies. For example, we can see that in Example 3 $C_1$ and $C_2$ are symmetric, so if there is an equilibrium in which the firm only exerts high effort for $C_1$ there is also an equilibrium in which the firm only exerts high effort for $C_2$.

Note that $\bar{t} = \underline{t}$ is not in general an appealing requirement for a parameterisation; one interpretation of the proposition would be that equilibria are generically unique when

[10]Note if the highest in-degree $d_{max}$ on network $\mathcal{N}$ satisfies $d_{max} \leq \underline{t}$, no equilibrium features high effort; if there is no customer on the network who can ever observe enough signals to be unwilling to purchase (for example, if $\pi$ is very large or $w$ is very small), then all customers always purchase and the firm will never exert high effort.

firm effort choice makes a negligible difference ($p_{11} \approx p_{11}$), which is not the most profound insight! For most interesting parameterisations of our model, multiple equilibria are likely to be possible. However, it worth noting that $\bar{t} = \underline{t} = 0$ is a potentially interesting case; this is the case where no matter what effort choice they expect, customers are turned off the firm by one bad review. Here it may be firm effort choice is still very productive, but customers have sufficiently low priors they avoid firms with even one bad review.

**Persistent High Effort:** Our results so far have demonstrated when it is possible in equilibrium for all customers beyond a certain date to copy their predecessors, and when we should expect to find unique equilibria. Given the pervasiveness of multiple equilibria in our model, without a strong equilibrium selection argument it is hard to make definitive predictions in general settings. However, we can provide insights into inference problems. In particular, we can suggest an answer to the following question: if we observe a firm receive persistent good reviews over time, what does this say about customer networks?

We will do so by comparing the maximum high effort that can be incentivised in equilibrium for two commonly studied network specifications: the complete network (where all customers observe all their predecessors) and the line network (where each customer only observes their immediate predecessor).

When the social network is complete and customers visit sequentially, $C_t$ visits in period $t$ and observes all the customers $C_{t-1}$ observed[11]. We can support $e_t = 1\ \forall t \leq t^*$ if $s_{t-s} = 0$, $e_t = 0$ otherwise, as an equilibrium strategy for the firm if the following conditions hold:

$$\begin{aligned}
w \geq& \frac{c(1-\delta)}{(1-p_{11})^{t^*-1}(p_{11}-p_{01})\delta^{t^*-1}(1-\delta^{T-t^*+1})} \\
w + u_0 \geq& \frac{(1-p_{11})^{N-1}\pi}{(1-p_{11})^{N-1}\pi + (1-\pi)} p_{01} \\
t^* =& \max_t \{t | t \in \mathcal{N}^+, t < N, \frac{(1-p_{11})^t\pi}{(1-p_{11})^t\pi + (1-\pi)} p_{11} \geq w + u_0\}
\end{aligned}$$

As $c \to 0$, it follows that the maximum number of customers who can receive high effort on a complete network is pinned down by:

$$t^* = \max_t \{t | t \in \mathcal{N}^+, t < N, \frac{(1-p_{11})^t\pi}{(1-p_{11})^t\pi + (1-\pi)} p_{11} \geq w + u_0\}$$

When the social network is a line network and $C_t$ visits in period $t$, such that $N(C_t) = \{C_{t-1}, C_{t+1}\}$ for $1 < t < N$, we know that there exist equilibria featuring herding, where

[11] A common approach to restricting customer observations in the reputations literature is to limit customers to visiting sequentially but observing only the outcomes of the past $K$ periods for finite $K$. In our setting, we note that either the finite memory setting has an equivalent equilibrium to a complete network with sequential visits, if for a complete network there is a critical customer among the first $K$, or in equilibrium features no high effort.

the firm exerts high effort only for $C_1$ (as in Example 1). However, if $\frac{\pi(1-p_{11})p_{01}}{\pi(1-p_{11}+(1-\pi)} < w + u_0 < \frac{\pi(1-p_{11})p_{11}}{\pi(1-p_{11}+(1-\pi)}$, there also exist other equilibria. In particular, suppose $N$ is even. Then for $c$ sufficiently small, there exist equilibria in which the firm sets $e_i = 1$ and customer always sets $a_i = 1$ if $i$ is odd, while if $i$ is even the firm sets $e_i = 0$ and $C_i$ sets $a_i = 1$ iff $s_{i-1} = 1$. In other words, odd customers always purchase and receive high effort, and even customers purchase if their immediate predecessor wrote a good review and receive low effort. This constitutes an equilibrium because if odd customers expect high effort they are happy to purchase at any information set and as such will ignore whether their predecessor purchases, while if even customers know they will receive low effort, they are only willing to purchase if they see a good review. As such, as $c \to 0$, if the total number of customers is even, the maximum number of customers who can receive high effort when customers only observe the previous period is $\frac{N}{2}$.

The following result is immediate:

**Proposition 3.** *If $N > \bar{N}$, the maximum number of customers who can receive high effort in equilibrium is higher with a line network than a complete network, for:*

$$2\lfloor ln\left(\frac{(1-\pi)(w+u_0)}{\pi(p_{11}-w-u_0)}\right\rfloor = \bar{N}$$

This result tells us that if the number of customers is sufficiently large, the maximum number of customers who can receive high effort in equilibrium is larger when customers observe only the previous period than when they observe the full history.

Note that this is the *maximum* number, not the expected number. For the line network the expected number of customers who receive high effort in equilibrium is also the maximum number, but for the complete network, the expected number is significantly lower, since if the firm receives a good review from $C_1$, they have revealed themselves to be high-skill and can quit high effort for all remaining customers. As a result, if we consider the maximum expected effort in equilibrium, the complete network compares even less favourably to the line network.

Our result suggests we should draw the following inference:

**Observation 1.** *Firms who receive persistent good reviews over time serve networks with intermediate customer connectivity.*

Note whilst our result uses the simple comparison between a line network and a complete network when effort is cheap, analogous results hold for higher effort costs[12] and observation structures where customers observe something between the full history and

[12] In particular, suppose for the line network, the firm will only exert high effort if it determines the next $M$ customers' decisions. Then we can find equilibria in which $C_1$ receives high effort, $C_2$ to $C_{M+1}$ herd (only purchasing if $s_1 = 1$), $C_{M+2}$ always purchases and receives high effort, $C_{M+3}$ to $C_{2M+3}$ herd, and so on. The same insight obtains; for $N$ sufficiently large, the highest effort possible in equilibrium will be higher for the line network than complete network.

their immediate predecessor. The driving force behind this result is that when all customers can observe sufficiently many previous periods, then beyond some date customers will only be willing to purchase if their is a good review in the history (regardless of what effort the customers expect from the firm). It follows that after that date, if a customer purchases all customers must believe the firm to be a high type for sure, and the firm will have no more incentives for effort. In contrast, if customers only observe a limited subset of the firm's past reviews, it is possible in equilibrium for the firm to restart high effort even if it has already achieved one good review, if not all later customers have learned about that good review. While our argument here is stated for fully revealing good news, the intuition is much more general; if customers read many reviews, their beliefs are harder to influence than when they each only read a few reviews. As such, the maximum high effort possible in equilibrium will be higher with information structures when customers do not have access to too many previous reviews, because for those structures incentives are more persistent.

# 4 Efficient Networks

In many settings we would like to understand which customer social networks are optimal from the perspective of various welfare measures. Social planners may sometimes have the ability to influence which social networks form: for example, a university building a new campus may have discretion over the layout of halls of residence and the number of communal spaces available. Policy makers and platforms may also have access to interventions which change the visibility of online customer reviews; for example, a reviews platform could add an auto-translate feature for foreign language reviews. In this section, we begin by presenting some simplifying results about what outcomes a planner designing customer network and order can implement in equilibrium. We then explore in greater depth the problem of a planner whose specific aim is to maximise total expected surplus. We also discuss properties of consumer and firm optimal networks, and under what circumstances these will differ from socially efficient networks.

We will assume that effort is ex-ante desirable from the perspective of total surplus (equivalently, we assume if effort were contractible high-skill firms would want to write contracts which induced high effort). Formally:

**Assumption 2.** $p_{11} - p_{01} > c$

This assumption best fits our motivations; exploring settings where reputational incentives can overcome moral hazard problems[13] in the absence of formal contracting.

[13]Parametrisations with inefficient effort, where high effort plays the role of noisy inefficient Spencian signalling, are also of theoretical interest, but are arguably less compelling as applications. As such, we assume effort is ex-ante efficient outright, for the sake of minimising caveats in stating our results.

We will continue to refer to a network $\mathcal{N}$ as a social network, but we reiterate that formally, $\mathcal{N}$ only represents an observation structure, and the reader is free to assign it different interpretations. When discussing a planner designing $\mathcal{N}$, for example, it may be appealing to interpret $\mathcal{N}$ as determining review visibility on an online platform.

We will consider the problem of a planner who before the first period publicly chooses a network-customer order pair, and suggests an equilibrium strategy profile to all players. We begin by providing some simplifying results which are independent of the planner's objective, before exploring the problem of a planner who seeks to maximise total expected surplus.

## 4.1 Feasible Outcomes

In order to understand optimal networks from the perspective of a specific objective, it is useful to build some insight into what outcomes it is feasible for a planner to implement by choosing a network-order pair.

**Bipartite Networks** We begin with a significant simplifying result; with fully revealing good news, any equilibrium outcome that can be supported on a network $\mathcal{N}$ can also be supported on a bipartite network $\mathcal{N}_B$. Formally, a network $\mathcal{N}$ with node set $\mathcal{C}$ and edge set $\mathcal{E}$ is a bipartite network if there exists $\mathcal{C}_1$ and $\mathcal{C}_2$ such that every edge in $\mathcal{E}$ connects a node in $\mathcal{C}_1$ to a node in $\mathcal{C}_2$.

**Proposition 4.** *If $\mathcal{N}, \succeq_\alpha$ is a network-order pair such that equilibrium strategy profile $(\sigma, \boldsymbol{\rho})$ generates ex-ante distribution over histories $\Delta(H_t)$, there exists a bipartite network $\mathcal{N}_\mathcal{B}$ with node sets ${\mathcal{C}_\mathcal{B}}^1$ and ${\mathcal{C}_\mathcal{B}}^2$, such that for network-order pair $\mathcal{N}_\mathcal{B}, \succeq_\alpha$, there exists an equilibrium strategy profile $(\sigma_B, \boldsymbol{\rho}_B)$ which also generates ex-ante distribution over histories $\Delta(H_t)$, and in equilibrium all $C_i \in {\mathcal{C}_\mathcal{B}}^1$ always purchase on the equilibrium path.*

This proposition states that for any outcomes that can be supported in equilibrium with a non-bipartite network, we can find a bipartite network in which customers visit in the same order and the same outcomes can be supported in equilibrium. The implication for a social planner is that (independent of planner objective function) there is no loss to the planner from restricting their search to only considering bipartite networks.

The reason is as follows: firstly, we can note there is no benefit to allowing customers who 'always buy' to observe one another, as they purchase at every equilibrium path info set. Secondly, because good news is fully revealing, anything one critical customer could learn from observing another critical customer, they could also learn by observing that critical customer's friends instead. In other words, if Bob is a critical customer and Alice observes only Bob, Alice takes exactly the same decisions in each on-path history as if she only observed Bob's friends but did not observe Bob, because any time she sees Bob purchase, she knows one of his friends wrote a good review, so whether Bob writes a good review himself is irrelevant. As a result, for any non-bipartite network in which critical customers learn from one another, we can design a new network, cutting out the middle man, in which they take exactly the same decisions at every on-path history but each critical customer observes only 'always buys' (note that this may involve *adding* links to the non-bipartite network as well as removing them).

It is worth emphasising that our assumption that $p_{10} = p_{00} = 0$ is crucial for this result; if good news is not fully revealing, a planner would not be able to restrict attention to only bipartite networks, because it would still be possible for customers to learn from one another after the first piece of good news was observed. In other words, 'sometimes buys' could learn from other 'sometimes buys' if good news were not fully revealing.

When choosing network-order pairs, planners face a trade-off between inducing high effort from the firm and providing information to later customers. How the planner values the trade-off will depend upon their specific objective function, but regardless of their objective it is useful to understand what the planner can achieve at either extreme. In other words, if we think of the planner as implicitly choosing a payoff vector from a set of feasible equilibrium payoffs, it is useful to characterise the extreme points of that set. As such, it is illuminating to ask; what should a planner who wants to induce as much high effort as possible do? What should a planner who wants to minimise the probability of customer 'mistakes' (buying from bad firms or failing to buy from good firms) do?

**Maximal High Effort** Example 3 illustrated that if $p_{11}$ is sufficiently large (high effort reviews are sufficiently informative) then it may be easier to incentivise high effort for always buys when later customers observe fewer reviews. This suggests the following question: if customers visit at $t = 1$ and $t = 2$, what is the largest number of customers who can visit at $t = 1$ and receive high effort in equilibrium?

To answer this question, it is useful to state the following lemma on when customer incentives are strongest:

**Lemma 1.** *If all customers visit in $t = 1$ and $t = 2$, and all $t = 1$ customers receive high effort, aggregate expected benefit from effort for firms is maximised if $t = 2$ customers have degree $m^*$, where*

$$m^* =\in \left\{ \lfloor \frac{-1}{\ln(1-p_{11})} \rfloor, \lceil \frac{-1}{\ln(1-p_{11})} \rceil \right\}$$

We can reason that roughly, the planner will be able to induce the most high effort by designing the network so that each $t = 2$ customer makes the largest possible contribution to aggregate incentives (since the more each $t = 2$ customer contributes to incentives, the fewer $t = 2$ customers are necessary to satisfy $t = 1$ IC constraints). The planner faces the following trade off; $t = 2$ customers with low degrees provide strong incentives to a few $t = 1$ customers, while $t = 2$ customers with high degrees provide weak incentives for many $t = 1$ customers. If a $t = 2$ customer has degree 1, they appear in 1 IC constraint at $t = 1$ while if they have degree 2 they appear in 2 IC constraints at $t = 2$, etc. Fixing number of $t = 1$ and $t = 2$ customers and maximising aggregate incentives with respect to degree of $t = 2$ customers gives our lemma.

Note that in general may not be feasible to give $t = 2$ customers degree $m^*$ in an efficient network-order pair. For example, it may be that $m^* > n_1$, where $n_1$ is the number of $t = 1$ customers, or it may be that $t = 2$ customers with degree $m^*$ cannot be symmetrically and evenly allocated between $t = 1$ customers without divisibility issues. In settings in which $\pi p_{01} - w \geq u_0$, it might also be the case that $t = 2$ customers with degree $m^*$ are not critical customers.

We can use our lemma to give an upper bound on the number of customers who can receive high effort in equilibrium. We do so using the following argument: if, using the strongest possible incentives (which may not be feasible due to divisibility issues), we would need more than $a$ $t = 2$ customers for each $t = 1$ customer, then on any *feasible* network $\frac{n_1}{n_2} < \frac{1}{a}$::

**Proposition 5.** *In a bipartite network featuring $n_1$ customers visiting at $t = 1$ and receiving high effort, and $n_2$ customers visiting at $t = 2$,, $\frac{n_1}{n_2} < \frac{1}{a}$, where $a$ is given by:*

$$a = \lfloor -\frac{c(1 - p_{11})(\ln[(1 - p_{11})^{\frac{1-p_{11}}{\ln(1-p_{11})}}])}{\delta(p_{11} - p_{01})w} \rfloor$$

Note that we bound by looking at the number of units of strongest possible incentives we would need to incentivise high effort and then rounding down. The logic here is the following; suppose that using the strongest incentives conceivable (which may not be feasible), we cannot satisfy aggregate IC when we have 3 $t = 2$ customers for each $t = 1$ customer. Then for any feasible network (featuring weakly weaker incentives), we must have at least 3 $t = 2$ customers for each $t = 1$ customer[14].

**Minimal Customer Mistakes** A planner who wants to minimise customer 'mistakes' can achieve arbitrarily few mistakes with a large enough number of customers if customers are always willing to purchase under their priors. In particular, the planner can do the following; allocate $n_1$ customers to $t = 1$ and $n_2$ customers to $t = 2$, and allow all $n_2$ customers to observe all $n_1$ customers. As $N = n_1 + n_2 \to \infty$, if the planner lets $n_1 \to \infty$ and $n_2 \to \infty$ but $\frac{n_1}{n_2} \to 0$, all $t = 2$ customers will have arbitrarily accurate posteriors and the proportion of customers who only purchase when $\theta = 1$ will go to 1 (this insight is from Acemoglu et al. (2011), who state it in a social learning setting). Note that since

[14]Note it does not follow that we should have 4 $t = 2$ customers for each $t = 1$ customers; if we only just fail to satisfy aggregate IC with 3 $t = 2$ customers for each $t = 1$, adding a single additional $t = 2$ customer who observes all $t = 1$ customers could be sufficient to motivate effort for all $t = 1$ customers/

$t = 2$ customers have arbitrarily accurate posteriors, the firm would never find high effort optimal at $t = 1$, since the expected change in revenue as a result of high effort becomes vanishingly small as $n_1 \to \infty$.

The problem is more interesting with a small number of customers when the planner considers choosing network-order pairs which do induce some high effort. Then the planner must contend with an additional effect; reviews from customers who receive high effort are more informative than reviews from customers who receive low effort. As a result, when the total number of customers is small, the planner may prefer to design a network in which $t = 2$ customers do not observe all $t = 1$ customers (as illustrated in example 3). Given the discrete nature of our model, it is hard to give a clean characterisation of what a planner who wants to minimise customer 'mistakes' should choose for a given number of customers, but if $p_{11} - p_{01}$ is sufficiently large (and $c$ sufficiently small) we would expect that with a small number of customers, the planner would do best choosing a network in which the firm is 'just' incentivised for high effort for some $t = 1$ customers, with the number of reviews later customers read limited by the firm's IC constraint for $t = 1$ customers.

## 4.2 Socially Optimal Networks

We now explore in more depth the problem of a planner who designs a network-order pair to maximise expected social surplus. Formally, we define total ex-ante expected surplus $S$ for a given network-order pair $(\mathcal{N}, \succeq_\alpha)$ and strategy profile $(\boldsymbol{\rho}, \sigma)$ as:

$$S = \sum_{t=1}^{T} \delta^{t-1} \left[ \sum_{C_i^t \in \mathcal{C}^t} E(a_i(s_i - ce_i) + (1 - a_i)u_0 | \boldsymbol{\rho}, \sigma, \mathcal{N}, \succeq_\alpha) \right]$$

In other words, it is the discounted sum of consumer and producer surplus. Note that total expected welfare $S$ is *not* a function of the price $w$, as purchases are transfers between customers and firms, and so cancel.

We say that a network-order pair $(\mathcal{N}^*, \succeq_\alpha^*)$ is *efficient* if there exists a strategy profile $(\boldsymbol{\rho}^*, \sigma^*)$ such that $(\boldsymbol{\rho}, \sigma)$ constitutes a PBE given $(\mathcal{N}^*, \succeq_\alpha^*)$, and $(\boldsymbol{\rho}^*, \sigma^*, \mathcal{N}^*, \succeq_\alpha^*) \in \operatorname{argmax}_{\boldsymbol{\rho}, \sigma, \mathcal{N}, \succeq_\alpha} S$. Note that efficient network-order pairs will not be unique in our setting; networks may have links that are irrelevant for equilibrium outcomes for a given order (say, if they link two customers who visit simultaneously), and as all customers in our model are ex-ante identical we can always swap $C_i$ and $C_j$ in $\mathcal{N}^*, \succeq_\alpha^*$ and the new network-order pair will also be efficient. Note also that as our setting can feature multiple equilibria (because customer willingness to pay depends on conjectures about firm effort), an efficient network-order pair $(\mathcal{N}^*, \succeq_\alpha^*)$ may also have other PBE strategy profiles $(\boldsymbol{\rho}', \sigma')$ such that $(\boldsymbol{\rho}', \sigma', \mathcal{N}^*, \succeq_\alpha^*) \notin \operatorname{argmax}_{\boldsymbol{\rho}, \sigma, \mathcal{N}, \succeq_\alpha} S$. If network-order pair $(\mathcal{N}^*, \succeq_\alpha^*)$ is *efficient*, we will refer to it as solving the social planner's problem.

What outcomes a social planner would like to induce will depend upon customer outside options. If $u_0 \leq 0$ trade between a low-skill firm and a customer is ex-ante desirable for total surplus (eating at a bad restaurant is still better than cooking at home), even if customers would prefer to avoid low-skill firms (because from their perspective, low skill firms are overpriced). If $u_0 > 0$ trade between a low-skill firm and a customer is ex-ante undesirable, and so a planner would like customers to avoid purchasing from low-skill firms.

Note this means that a planner interested only in maximising total surplus $S$ has potentially very different priorities to a planner interested in maximising consumer surplus. In particular, a planner interested in total surplus does not care how much a customer pays for a firm's services, only if the customer's *gross* benefit from purchasing is positive or negative. Phrased another way: if Firm A and Bob both have linear payoffs in wealth, Firm A values an apple at £0.00 and Bob values an apple at £0.01, Bob purchasing an apple from Firm A at price £1000 maximises total surplus, even if it is very harmful for Bob's consumer surplus! As a result, it is potentially very important for policy makers to be clear about whether they are interested in total surplus or consumer surplus when considering policy interventions. We discuss how consumer optimal network-order pairs relate to socially optimal network-order pairs at the end of this section.

The decision that customers would take under their prior plays an important role in shaping the planner's problem. If $\pi p_{01} - w \geq u_0 > -w$, uninformed customers (inheriting prior) are willing to purchase regardless of their conjectures about firm effort. Hence if no customers learn anything about the firm (for example, if the social network is empty), all customers will purchase. Customers will only stop purchasing from the firm if they have seen a sufficient amount of bad news and revised their beliefs down relative to their prior. The planner can guarantee surplus of $S_0 = N\pi p_{01}$ by allocating all customers to $t = 1$ and choosing any network.

In contrast if $\pi p_{01} - w < u_0 \leq \pi p_{11} - w$, uninformed customers will purchase only if they expect high effort. Hence for any trade to occur a planner will need to incentivise high effort for a first wave of customers, from whom later customers can learn about firm type. The planner can guarantee surplus of $S_\emptyset = Nu_0$ by allocating all customers to $t = 1$ and choosing any network.

We can significantly simplify the customer orders a planner needs to consider for optimal network-order pairs.

**Proposition 6.** *If $\delta < 1$ and the optimal network $\mathcal{N}^*$ is a complete bipartite network, then optimal order $\succeq^*$ allocates all customers to periods $t = 1$, $t = 2$, or $t = 3$. If $\pi p_{01} - w < u_0$, any optimal order $\succeq^*$ allocates all customers to periods $t = 1$ or $t = 2$.*

This proposition states that if the planner discounts the future, and all 'sometimes buys' observe all 'always buys', the planner will allocate customers to at most the first three periods. If customers are only willing to purchase expecting high effort under their priors, the planner will allocate all customers to at most the first two periods[15]. The role of assuming $\delta < 1$ here is simply to break ties; if $\delta = 1$ then we will have many optimal orders since the planner is happy to leave arbitrary gaps between customers (provided during the gaps customers learn nothing new).

On an intuitive level, we can think of this result as saying we can divide the optimal equilibrium path into up to three phases: reputation decay for the firm, high effort for the firm, and informed decisions for customers. This result claims that each of these phases can be achieved in one period if the network is a complete bipartite network.

To sketch the argument, firstly note that for an impatient planner, it is optimal for all critical customers visit at once. Proposition 4 tells us there is no benefit to critical customers observing one another, so with $\delta < 1$ the planner will do best when they visit simultaneously.

Secondly, note that expected total effort is higher when 'always buys' who receive high effort at some histories visit simultaneously. If all critical customers observe all 'always buys' (as on the complete bipartite network), then if 'always buys' visit sequentially the firm can quit high effort [16] after their first good review. Hence moving later 'always buys' earlier to visit at the same time as the first customer who sometimes receives high effort increases expected effort for the moved customers while leaving expected effort for other customers unchanged. In a symmetric setting[17], removing the firm's ability to wait-and-see with effort choices can only strengthen aggregate incentives.

Finally, we know from example 2 that it may sometimes not be possible to incentivise high effort unless the firm has already received sufficiently many bad reviews, if customers have high priors. As a result, for some parameterisations the optimal order may involve an initial phase of reputation decay, to incentivise effort at a later date. If the planner has $\delta < 1$ the planner will do best allocating all of the 'always buys' who receive low effort to a single period.

Our results so far tell us a planner can restrict their search to bipartite networks, and which orders a planner should consider for a *complete* bipartite network. We now give conditions for when a *complete* bipartite network is optimal, and argue that if the number

[15]Note there is no obligation on the planner to use more than one period: if the planner wants to induce full trade or no trade equilibria, with $\delta < 1$ they will do best allocating all customers to $t = 1$.

[16]Note if the planner maximised firm profits (or if effort were ex-ante inefficient and had purely signalling value) this would be a beneficial effect; with inefficient high effort the planner would like to let the firm wait and see specifically to allow them to quit effort after the first success.

[17]Asymmetric settings are more complicated; it may be that if $C_j$ has a lower degree than all the other 'always buys', they cannot receive high effort if they visit at the same time as other 'always buys', but can receive high effort in some histories if they visit later. In practice, optimal networks are likely to be approximately symmetric (the planner has no intrinsic reason to provide a customer with stronger incentives than necessary, and for critical customers there is diminishing marginal value to observing an extra review, both of which point the planner towards symmetry), but given the discreteness of our setting, we cannot rule out asymmetry due to divisibility issues.

of customers is sufficiently large and high effort is sufficiently productive, the planner will not want to choose a complete bipartite network.

**Optimality of Complete Bipartite Network** To help understand when a planner would like to use a complete bipartite network, we can ask: when are learning and incentives are complements for a planner? In other words, when does a more connected network improve both customer decision making and firm incentives for effort?

Recall in lemma 1 we found:

$$m^* = \left\{ \lfloor \frac{-1}{\ln(1-p_{11})} \rfloor, \lceil \frac{-1}{\ln(1-p_{11})} \rceil \right\}$$

where $m^*$ is the degree for $t=2$ customers which maximises their contribution to aggregate incentives for $t=1$ customers. Understanding when incentives are strongest, we can note the following; if we fix the node sets of a bipartite network, and all critical customers currently have a degree below $m^*$, increasing all of their degrees by one should strengthen aggregate incentives *and* improve customer learning. Hence the following result holds:

**Proposition 7.** *In the class of bipartite networks with node sets* ${\mathcal{C}_\mathcal{B}}^1$ *and* ${\mathcal{C}_\mathcal{B}}^2$ *and* $|{\mathcal{C}_\mathcal{B}}^1| = n_1$, $|{\mathcal{C}_\mathcal{B}}^2| = n_2$, *where all members of* ${\mathcal{C}_\mathcal{B}}^1$ *visit the firm at* $t=1$, *the complete bipartite network is optimal if* $n_1 \leq m^*$ *and* $u_0 > 0$.

In other words, if the highest possible number of 'always buys' it is possible for a critical customer to observe is below $m^*$, it is socially optimal to let all critical customers observe all 'always buys'.

In contrast if $n_1 > m^*$, the planner may face a trade off between learning and incentives; other things equal, a (uniformly) more connected network provides weaker incentives even if it may provide better learning. The driving force behind this trade off is simple; as customers have access to more information, the chance that any one review changes their decision goes to zero. Hence, if customers read too many reviews each, a firm may not find high effort worthwhile when serving them, reasoning that they are likely to learn the firm's type from the other reviews regardless.

Note that the condition $n_1 \leq m^*$ may be very hard to satisfy; for example, if $p_{11} = 0.6$, we have $m^* = 2$, so with more than two customers who always buy, increasing the number of reviews later customers can read harms firm incentives. We can also note $n_1 \leq m^*$ is not a necessary condition for a complete bipartite network to be optimal; if we fix the two node sets of a bipartite network, the planner may prefer to use the weakest incentives that still motivate a firm to effort. If we have customers visiting at $t=1$ and $t=2$, the planner can improve social welfare by adding links between $t=1$ and $t=2$ customers provided this does not affect which customers a firm wants to exert high effort for.

Observing that if we increase the number of customers beyond a certain point, increasing connectivity further will weaken firm incentives, we can say:

**Proposition 8.** *There exists some $n'$ and some $\beta \geq 0$ such that if $N > n'$ and $p_{11} - p_{01} - c \geq \beta$, if high effort is possible in equilibrium for some network-order pair, no complete bipartite network is an efficient network.*

This result constitutes two claims: firstly, if high effort is sufficiently productive, as $N$ becomes large a social planner will want to incentivise high effort for more than one customer; secondly, as $N$ becomes large, it will become impossible to incentivise high effort for more than one customer in equilibrium.

This is true because incentives for firm effort are strongest when the watching customers have finite degree. For example, consider the case where $p_{11} = 1$; then it is impossible for more than one customer to receive high effort on a complete bipartite network in which customers visit at $t = 1$ or $t = 2$. However, for $n$ sufficiently large, it will be possible to incentivise high effort for more than one customer on a disjoint network, if each 'always buy' customer has enough exclusive followers. Where a planner places significantly more weight on incentives than customer learning, non-complete social networks may prove optimal.

A corollary with important policy implications is:

**Corollary 2.** *Efficient networks may feature disjoint components.*

This tells us that interventions such as auto-translating foreign language reviews may be harmful for social welfare, if the result is that they connect two previously disjoint cliques on which a firm previously had to build a reputation separately.

Note however that because the planner is risk neutral, there is no inherent power in disjoint networks over connected networks with similar properties. For example, if we divide customers into equal numbers of 'always buys' and critical customers, any 2-regular bipartite network will give the same expected surplus, whether it is a circle network or divides the network into disjoint cliques of 4 customers each. As such, our insight is that disjoint networks may be socially preferable to complete networks, *not* that disjoint networks are uniquely socially optimal.

**Customer Optimal Networks:** We have focused on socially optimal networks, but can argue that, for many parameterisations, consumer optimal networks will be qualitatively very similar to socially optimal networks. In particular, when $(p_{11} - p_{01})w > c$ and $u_0 > 0$, both total surplus and consumer surplus satisfy the following properties: they are linear and increasing in $e_i$ and $P(a_i|\theta = 1)$, and linear and decreasing in $P(a_i|\theta = 0)$, for all $C_i$. In other words, when effort is ex-ante efficient and customers' outside options generate more surplus than purchasing from a bad firm, both total and consumer surplus agree high effort is good, purchasing from high-skill firms is good, and purchasing from low-skill firms is bad.

The exact rate at which a planner values the trade-off between firm effort and customer learning will depend on their specific objective; for example, we know the objective of a planner who cares about total surplus does not depend on prices, while a planner who cares about consumer surplus values customers avoiding bad firms more when prices are higher. As such, for a given parameterisation, the socially optimal network is unlikely to be exactly the customer optimal network. Nonetheless, we can note the following; when $(p_{11} - p_{01})w > c$ and $u_0 > 0$, consumer optimal networks will be qualitatively similar to socially optimal networks (if not quantitatively identical), because consumer and social surplus make qualitatively similar trade-offs between firm effort and learning. As such, we can develop an understanding of the properties of consumer optimal networks by studying socially optimal networks; there is unlikely to be significant benefit to a formal analysis of customer optimal network-order pairs on top of an analysis of socially optimal network-order pairs[18].

**Firm Optimal Networks:** When $w < \pi p_{01}$, the question of firm optimal networks is simple; if customers are always willing to purchase from an unknown firm under their prior, regardless of firm type the firm is best off with a network order-pair which induces no learning and no incentives for effort (for example, any network-order pair where all customers visit at $t = 1$). In contrast, if $w > \pi p_{01}$, uninformed customers must expect to receive high effort to be willing to purchase, so (regardless of firm type) the optimal network-order pair must be one which would induce some high effort from a high-skill firm.

We can see that when $w > \pi p_{01}$ the firm optimal network will depend upon firm type. If the firm is a high-skill type, they value both incentivising their own effort (so that customers are willing to purchase under their prior) and providing later customers with the best information possible (as if later customers learn the firm is high-skill they will visit). In contrast, if the firm is a low-skill type, while they want the network to be one which would incentivise effort from a high-skill firm (so that customers are willing to purchase under their prior), they do not care about the information quality of later critical customers, since with fully revealing good reviews, no critical customer will ever purchase from a low-skill firm, as a low-skill firm cannot receive good reviews.

We can reason that a low-skill firm should prefer a network-order pair which maximises aggregate incentives, as discussed in proposition 5. In contrast, it is possible a high-skill firm may not want to maximise aggregate incentives; they may prefer ex-ante to have fewer $t = 1$ customers and more $t = 2$ customers (even though $t = 2$ customers do not always purchase), if they need to exert high effort for all $t = 1$ customers to motivate

[18]Phrased another way, we can conjecture the following: if for some parameterisation $\mathcal{N}^*, \succeq^*$ is a consumer optimal network, there exists a reparameterization of the model for which $\mathcal{N}^*, \succeq^*$ is socially optimal. In other words, we do not expect consumer optimal networks to explore parts of the network-order space not explored by socially optimal networks.

them to visit. This is particularly clear if $p_{11}$ is very close to 1; if $(p_{11} - p_{01})\delta w > c$, a high-skill firm prefers to have one $t = 1$ customer, observed by $N - 1$ $t = 2$ customers, while with $N = 2n$, the low-skill firm prefers a one-to-one matching between $n$ $t = 1$ and $n$ $t = 2$ customers (each reviewer is viewed by one reader, each reader reads one review). As such, there is a sense in which a low-skill firm's ex-ante preferences over networks may appear closer to consumer or a social planner's preferences than a high-skill firm's[19].

# 5 Conclusions

We have shown that when incentives are reputational, customer social networks matter for incentives; a firm has incentives for effort only if the quality of service they provide will be observed by a customer whose purchase decision depends upon reviews. The strength of these incentives depends on the network structure; incentives are stronger when many critical customers (and associated herds) are watching, but weaker when those critical customers have access to many other reviews. We have argued that as service generated by high effort becomes fully informative about firm type, efficient social networks can feature disjoint components if high effort is sufficiently desirable.

For tractability, we have focused on a simple model in which all players have perfect information about the network and customer order, and in which good service fully reveals the firm's type. Many of our insights generalise naturally to broader settings. For the remainder of this section, we discuss possible generalisations, and relate our results to alternative approaches to modelling reputation.

We would expect herding to be a feature of equilibrium which carried over to more general settings. In particular, wherever customers make a binary decision between purchasing and not purchasing at a fixed price, observing a customer's purchase decision will reveal information about their beliefs, and so herding is always possible when a poorly informed customer observes a well informed customer. If customer $C_i$ observes $C_j$, and believes that customer $C_j$ is sufficiently better informed than them (for example, if customer $C_i$'s neighbours are a subset of customer $C_j$'s neighbours), then customer $C_i$ should purchase if $C_j$ does. This would be true even if we allowed low skill firms to sometimes generate good service, or made quality of service a continuous random variable.

Many of our results generalise naturally to settings with an infinite population of customers. For example, with one customer per period our results for directed rooted tree and complete networks generalise easily to infinite networks.

We might be interested in settings featuring weaker assumptions on the firm's knowledge of the network. If a firm is uncertain about who observes whom (either because they

[19]Of course, conditional on the firm being a low type, customers and social planners would prefer a network which induced no trade; the point here is that when firms have discretion over observation structures, customers should be more sceptical of firms who choose to implement observation structures which appear ex-ante customer optimal.

do not know the network structure or because they do not know where on the network a customer is located), then for each customer the firm will assign probabilities to that customer being observed by each other customer. Each customer will have a cut-off strategy in beliefs, which will depend on their location on the network[20]. We would expect analogues of many of our insights to hold in this more general setting; the firm will exert high effort for a customer if they believe it is *sufficiently likely* that customer is observed by a critical customer. Similarly, if a customer believes it is sufficiently likely that their friend is well informed, they will interpret that friend purchasing from the firm as a sign the firm is high quality and choose to purchase themselves.

While we can speculate about how a firm would behave if they did not know the network or order of customers, a full analysis is not easy. Note that the firm's problem is potentially very complicated; their strategy in a given period will depend on in how many previous periods they achieved high effort and in how many previous periods they were rejected. Historic rejections may suggest that herding has already begun to the firm's detriment, meaning the firm may be deterred from future high effort, believing there are few customers remaining whose purchase decisions they can influence.

We might also want to explore endogenous customer orders. In many settings it is unrealistic to assume that the orders in which customers visit is exogenously fixed; for example, customers generally choose when to book to visit a new restaurant. We would expect the order in which customers visit to depend upon their location on the social network; for example, central customers might visit early, uncertain about the firm's type but knowing the firm has an incentive for high effort if skilled, while less central customers might choose to visit later, letting their better connected friends act as reviewers first. This is work in progress.

It is worth noting that in the presence of reputational incentives, the predictions of a model will depend importantly upon the possible types of a firm. We have considered a model in which a high-skill strategic firm seeks to distinguish itself from a low-skill commitment type firm which always provides bad service. As a result, incentives are driven by a desire by the firm to distinguish themselves; the firm has an incentive to exert high effort if doing so can statistically identify them as high-skill rather than low skill to later customers. A consequence of this is that incentives strengthen over time in 'bad histories'; a firm who is yet to distinguish themselves as high-skill has a finite number of opportunities to do so in equilibrium, and finds incentives strengthen as the number of remaining opportunities reduces (other things equal).

In contrast, suppose we considered a model in which a low-skill strategic firm tried to imitate a high-skill commitment type (alternatively, a strategic firm seeking to imitate a

[20]If we assumed the firm knew topological properties of the network but not the sequence of customers, an infinite network would simplify the firm's problem significantly; with an infinite network known to possess some kind of regular topology, the firm will not need to worry about the correlation in customer 'types' that occurs when drawing nodes from a finite network without replacement.

Stackelberg type). Here, the firm has incentives for effort if doing so makes it harder for customers to identify them as low-skill. The consequence of this is that while incentives may strengthen in 'good histories' (as the chance of successfully going undetected given high effort goes up over time), they weaken in 'bad histories', as a firm who has already been 'caught' has nothing to gain from further high effort. In other words, if good firms seeks to distinguish themselves from bad, firms with good reviews are less likely to be exerting high effort; if bad firms seek to imitate good, firms with good reviews are more likely to be exerting high effort.

An important distinction between our model and the general literature on reputations is that the literature focuses on settings in which the long-run player actions are statistically identified by observing an unbounded history of play. A key component to this assumption is that regardless of short-run player actions, outcomes are informative about long run player actions. In our setting, this would be equivalent to requiring that observing the quality of service of an earlier friend is informative about firm effort even if that earlier friend did not purchase (for example, if firm effort generated externalities and so affected payoff from outside option). If long-run player actions are statistically identified even in 'bad histories' (in which almost all customers do not purchase), later customers can still learn about firm type if they observe the full history. In contrast, in our model, if almost all of a customer's earlier friends have not purchased, with positive probability that customer will hold incorrect beliefs about the firm's type regardless of their visit date. Hence, compared to the standard reputations literature, our results place much more importance on ensuring initial customers purchase; once customers stop purchasing in our setting a firm can earn no more revenue, whereas if customers can learn about a firm's type even if their friends do not purchase, firms understand they can recover from initial no- purchases and a bad reputation is always salvageable for a good firm.

# References


Acemoglu, Daron et al. (2011). “Bayesian Learning in Social Networks”. In: *The Review of Economic Studies* 78.4, pp. 1201–1236. ISSN: 00346527, 1467937X. URL: http://www.jstor.org/stable/41407059.

Banerjee, Abhijit V. (1992). “A Simple Model of Herd Behavior*”. In: *The Quarterly Journal of Economics* 107.3, pp. 797–817. ISSN: 0033-5533. DOI: 10.2307/2118364. eprint: https://academic.oup.com/qje/article-pdf/107/3/797/5298496/107-3-797.pdf. URL: https://doi.org/10.2307/2118364.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades”. In: *Journal of Political Economy* 100.5, pp. 992–1026. ISSN: 00223808, 1537534X. URL: http://www.jstor.org/stable/2138632.

Board, Simon and Moritz Meyer-ter-Vehn (2021). “Learning Dynamics on Social Networks”. working paper.

Galeotti, Andrea and Sanjeev Goyal (2009). “Influencing the influencers: a theory of strategic diffusion”. In: *The RAND Journal of Economics* 40.3, pp. 509–532.

Golub, Ben and Evan Sadler (2016). “Learning in Social Networks”. In: *Oxford Handbook of Economic Networks*.

Mailath, George J. and Larry Samuelson (2001). “Who Wants a Good Reputation?” In: *The Review of Economic Studies* 68.2, pp. 415–441. ISSN: 0034-6527. DOI: 10.1111/1467-937X.00175. eprint: https://academic.oup.com/restud/article-pdf/68/2/415/4369473/68-2-415.pdf. URL: https://doi.org/10.1111/1467-937X.00175.

Pei, Harry (2021). “Reputation Building under Observational Learning”. working paper. URL: https://cpb-us-e1.wpmucdn.com/sites.northwestern.edu/dist/4/2519/files/2021/08/Main_text.pdf.

Sgroi, Daniel (2002). “Optimizing Information in the Herd: Guinea Pigs, Profits, and Welfare”. In: *Games and Economic Behavior* 39.1, pp. 137–166. URL: https://ideas.repec.org/a/eee/gamebe/v39y2002i1p137-166.html.