Personalised Rating

Umberto Grandi, James Stewart and Paolo Turrini

Abstract

We study a rating system in which a set of individuals (e.g., the customers of a restaurant) evaluate a given service (e.g, the restaurant), with their aggregated opinion determining the probability of all individuals to use the service and thus its generated revenue. We explicitly model the influence relation by a social network, with individuals being influenced by the evaluation of their trusted peers. On top of that we allow a malicious service provider (e.g., the restaurant owner) to bribe some individuals, i.e., to invest a part of his or her expected income to modify their opinion, therefore influencing his or her final gain. We analyse the effect of bribing strategies under various constraints, and we show under what conditions the system is bribery-proof, i.e., no bribing strategy yields a strictly positive expected gain to the service provider. We also study the computational complexity of bribery, showing that establishing the existence of an optimal manipulation for the attacker is **NP**-complete, even with full knowledge of the underlying network structure.[a]

[a]This paper summarises two publications that appeared in the proceedings of the *25th International Joint Conference on Artificial Intelligence* (Grandi and Turrini, 2016) and of the *32nd AAAI Conference on Artificial Intelligence* (Grandi et al., 2018).

1 Introduction

Imagine to be the owner of a new and still relatively unknown restaurant. The quality of food is not spectacular and the customers you have seen so far are only limited to a tiny number of friends of yours. Your account on Tripadvisor® has received no review and your financial prospects look grim at best. There is one easy solution to your problems: you ask your friends to write an enthusiastic review for you, in exchange for a free meal. After this, Tripadvisor® lists your restaurant as excellent and the number of customers, together with your profit, suddenly florishes.

Systems such as Tripadvisor®, where a small proportion of customers writes reviews and influences a large number of potential customers, are not *bribery-proof*: each restaurant owner - or the owner of whichever service - is able to offer a compensation - monetary or not - in exchange for positive evaluation, having an impact on the whole set of potential customers. Tripadvisor® is based on what we call "Objective Rating", or $\mathbb{O}$-rating: individual evaluations are aggregated into a single figure, which is seen by, and thus influences, every potential customer.

What we study in this paper is a system in which each individual *only* receives the evaluation given by the set of trusted peers, his or her friends, and only this aggregated opinion influences his or her decision. This is what we call "Personalised Rating", or $\mathbb{P}$-rating, which can be seen a generalisation of $\mathbb{O}$-rating in which influence has a complex network-structure. So, while in the case of $\mathbb{O}$-rating the restaurant owner knows exactly how influence flows among the customers, this might not be the case with $\mathbb{P}$-rating.

Our contribution. We analyse the effect of bribing strategies in the case of $\mathbb{O}$-rating and $\mathbb{P}$-rating under various constraints, depending on the presence of customers who do not express any opinion and the knowledge of the network by the service provider: the exact network is known, the network is known but not the customers' exact position, the network is completely unknown. We show under what conditions the system is bribery-proof, i.e., there is no bribe yielding a strictly positive expected gain to the service provider, and we provide algorithms for the computation of (all) optimal bribing strategies when they exist.

Intuitively, being able to know and bribe influential customers is crucial for guaranteeing a positive expected reward of a bribing strategy. However, while with large populations of non-voters "random" bribes can still be profitable, the effect of $\mathbb{P}$-rating is largely different from that of $\mathbb{O}$-rating and, as we show, the expected profit in the former can be severely limited and drops below zero in all networks, under certain conditions on the cost of bribes.

Moreover, we show that even if $\mathbb{P}$-rating is manipulable in theory, the problem of manipulating it is intractable in practice. In particular, we establish that even when the attacker has full knowledge of the network the problem of determining the existence of a manipulation strategy guaranteeing at least a given reward — and, notably, an optimal one — is **NP**-complete. We do so by giving a polynomial-time reduction from the problem of finding an independent set of a given size k in a 3-regular graph.

Paper structure Section 2 presents the basic setup, introducing $\mathbb{O}$-rating, $\mathbb{P}$-rating and bribing strategies. Section 3 focusses on $\mathbb{O}$-rating, studying its bribery-proofness under various knowledge conditions. Section 4 evaluates $\mathbb{P}$-rating against the same knowledge conditions. In Section 5 we compare the two systems, taking the cost of bribery into account. In Section 6 we establish the **NP**-completeness of computing an optimal bribing strategy and of deciding the possibility for successful manipulation. We conclude by summarising the main findings and pointing at future research directions (Section 7). Due to space constraints, some of the proofs can be found in appendix.

1.1 Related Work

Our approach relates to several research lines in artificial intelligence, game theory and (computational) social choice (Brandt et al., 2015).

Network-based voting and mechanism design. We study social networks in which individuals' local decisions can be manipulated to modify the resulting global properties. A similar approach is taken by Apt and Markakis (2014) and Simon and Apt (2015), which study the changes on a social network needed to make a certain product adopted among users. The work of Conitzer et al. (2010), Bu (2013), Todo and Conitzer (2013) and Brill et al. (2016) on the effect of adding fake profiles to a social network is also relevant in its being close to the problem of bribing. Finally, an extremely relevant line of research is the work of Alon et al. (2015) and Lev and Tennenholtz (2017), who looks at theoretical guarantees for group recommendations, as well as papers that have looked at social network-based recommendations, such as Andersen et al. (2008). See also the recent survey by Grandi (2017) for relevant literature on the interplay between mechanisms for collective choice and social networks.

Lobbying and Bribery. Our framework features an external agent trying to influence individual decisions to reach his or her private objectives. Lobbying in decision-making is an important problem in the area of social choice, from the seminal contribution of Helpman and Persson (1998) to more recent studies in multi-issue voting (Christian et al., 2007). Lobbying and bribery are also established concepts in computational social choice, with their computational complexity being analysed extensively (Faliszewski et al., 2009; Baumeister et al., 2011; Bredereck et al., 2014, 2016).

Reputation-based systems. We study the aggregation of possibly insincere individual evaluations by agents that can influence one another through trust relations. In this sense ours can be seen as a study of reputation in Multi Agent Systems, which has been an important concern of MAS for the past decades (Conte and Paolucci, 2002; Sabater and Sierra, 2005; Garcin et al., 2009). In particular, our framework treats reputation as a manipulable piece of information, not just a static aggregate of individual opinions, coherently with the work of Conte et al. (2008) and Pinyol and Sabater-Mir (2013).

2 Basic setup

In this section we provide the basic formal definitions of our setting.

2.1 Restaurant and customers

Our framework features an object r, called *restaurant*, being evaluated by a finite non-empty set of individuals $C = \{c_1, \ldots, c_n\}$, called *customers*. Customers are connected by an undirected graph $E \subseteq C \times C$, called the *customers network*. Given $c \in C$ we call $N(c) = \{x \in C \mid (c, x) \in E\}$ the *neighbourhood* of c. We assume that $c \in N(c)$ for all c.

Customers concurrently submit an *evaluation* of the restaurant, drawn from a set of values $Val \subseteq [0, 1]$, together with a distinguished element $\{*\}$, symbolising no opinion. Examples of values are the set $[0, 1]$ itself, or a discrete assignment of 1 to 5 stars, as common in online rating systems. We make the assumption that $\{0, 1\} \subseteq Val$ and that Val is closed under the operation $\min\{1, x + y\}$ for all $x, y \in Val$. The vast majority of known rating methods can be mapped onto the $[0, 1]$ interval and analysed within our framework.

We represent the evaluation of the customers as a function $eval : C \to Val \cup \{*\}$ and define $V \subseteq C$ as the subset of customers that expresses an evaluation over the restaurant, i.e., $V = \{c \in C \mid eval(c) \neq *\}$. We refer to this set as the set of *voters* and we assume it to be always non-empty, i.e., there is at least one customer that expresses an evaluation.

2.2 Two rating systems

In online rating systems such as Tripadvisor® every interested customer can see - and is therefore influenced by - (the average of) what the other customers have written. We call this method $\mathbb{O}$-$\mathtt{rating}$, which stands for *objective rating*.

Given an evaluation function *eval* of a restaurant, the associated $\mathbb{O}$-$\mathtt{rating}$ is defined as follows:

$$\mathbb{O}\text{-}\mathtt{rating}(eval) = \operatorname*{avg}_{c \in V} eval(c)$$

Where avg is the average function across real-valued $eval(c)$, disregarding $*$. We omit *eval* when clear from the context.

$\mathbb{O}$-$\mathtt{rating}$ flattens individual evaluations into a unique objective aggregate, the rating that a certain restaurant is given. What we propose is a refinement of $\mathbb{O}$-$\mathtt{rating}$, which takes the network of influence into account. In this system customers are *only* interested in the evaluation of other customers they can trust, e.g., their friends. We call our method $\mathbb{P}$-$\mathtt{rating}$, which stands for *personalised rating*. It is defined for a pair customer-evaluation $(c, eval)$ as follows:

$$\mathbb{P}\text{-}\mathtt{rating}(c, eval) = \operatorname*{avg}_{k \in N(c) \cap V} eval(k)$$

So the $\mathbb{P}$-$\mathtt{rating}(c, eval)$ calculates what customer c comes to think of the restaurant, taking the average of the opinions of the customers c is connected to. Again we omit *eval* whenever clear from the context.

Observe that in case a customer has no connection with a voter, then $\mathbb{P}$-$\mathtt{rating}$ is not defined. To facilitate the analysis we make the technical assumption that *each customer is connected to at least one voter*. Also observe that when $E = C \times C$, i.e., in case the network is complete and each individual is influenced by each other individual, then for all $c \in C$ and *eval* we have that $\mathbb{P}$-$\mathtt{rating}(c, eval) = \mathbb{O}$-$\mathtt{rating}(eval)$.

2.3 Utilities and strategies

We interpret a customer evaluation as a measure of his or her *propensity* to go to the restaurant. We therefore assume that the utility that a restaurant gets is proportional to its rating. To simplify the analysis we assume a factor 1 proportionality.

The case of $\mathbb{O}$-$\mathtt{rating}$. For the $\mathbb{O}$-$\mathtt{rating}$, we assume that the initial utility u^0 of the restaurant is defined as:

$$u^0_{\mathbb{O}} = |C|\mathbb{O}\text{-}\mathtt{rating}(eval).$$

Intuitively, the initial utility amounts to the number of customers that actually go to the restaurant, weighted with their (average) predisposition.

At the initial stage of the game, the restaurant owner receives u^0, and can then decide to invest a part of it to influence a subset of customers and improve upon the initial gain. We assume utility to be fully transferrable and, to facilitate the analysis, that such transfers translate directly into changes of customers' predispositions.

Definition 1. *A strategy is a function $\sigma : C \to Val$ such that $\sum_{c \in C} \sigma(c) \leq u^0$.*

Definition 1 imposes that strategies are *budget balanced*, i.e., restaurants can only pay with resources they have.

Let Σ be the set of all strategies. We denote σ^0 the strategy that assigns 0 to all customers and we call *bribing strategy* any strategy that is different from σ^0. After the execution of a bribing strategy, the evaluation is updated as follows:

Definition 2. *The evaluation $eval^\sigma(c)$ after execution of σ is $eval^\sigma(c) = \min\{1, eval(c) + \sigma(c)\}$, where $* + \sigma(c) = \sigma(c)$, if $\sigma(c) \neq 0$, and $* + \sigma(c) = *$, if $\sigma(c) = 0$.*

In this definition we are making the assumption that the effect of bribing a non-voter to vote is equivalent to that of bribing a voter that had a 0-level review, as, intuitively, the individual has no associated predisposition to go to the restaurant.

A strategy is called *efficient* if $\sigma(c) + eval(c) \leq 1$ for all $c \in C$. Let $B(\sigma) = \{c \in C \mid \sigma(c) \neq 0\}$ be the set of bribed customers. Let V^σ be the set of voters after the execution of σ. Executing σ induces the following change in utility:

$$u^\sigma_{\mathbb{O}} = |C|\mathbb{O}\text{-}\mathtt{rating}(eval^\sigma) - \sum_{c \in C} \sigma(c).$$

Intuitively, $u^\sigma_{\mathbb{O}}$ is obtained by adding to the initial utility of the restaurant the rating obtained as an effect of the money invested on each individual minus the amount of money spent.

We define the revenue of a strategy σ as the marginal utility obtained by executing it:

Definition 3. *Let σ be a strategy. The* revenue *of σ is defined as $\mathbf{r}_{\mathbb{O}}(\sigma) = u^\sigma_{\mathbb{O}} - u^0$. We say that σ is* profitable *if $\mathbf{r}_{\mathbb{O}}(\sigma) > 0$.*

Finally, we recall the standard notion of dominance:

Definition 4. *A strategy σ is* weakly dominant *if $u^{\sigma}_{\mathbb{O}} \geq u^{\sigma'}_{\mathbb{O}}$ for all $\sigma' \in \Sigma$. It is* strictly dominant *if $u^{\sigma}_{\mathbb{O}} > u^{\sigma'}_{\mathbb{O}}$ for all $\sigma \in \Sigma$.*

Hence a non-profitable strategy is never strictly dominant.

The case of $\mathbb{P}$-rating. The previous definitions can be adapted to the case of $\mathbb{P}$-rating as follows:

$$u^0_{\mathbb{P}} = \sum_{c \in C} \mathbb{P}\text{-rating}(c, eval)$$

which encodes the initial utility of each restaurant, and

$$u^{\sigma}_{\mathbb{P}} = \sum_{c \in C} \mathbb{P}\text{-rating}(c, eval^{\sigma}) - \sum_{c \in C} \sigma(c)$$

which encodes the utility change after the execution of a σ. Finally, let the revenue of σ be $\mathbf{r}_{\mathbb{P}}(\sigma) = u^{\sigma}_{\mathbb{P}} - u^0_{\mathbb{P}}$. If clear from the context, we use $\mathbb{P}\text{-rating}^{\sigma}(c)$ for $\mathbb{P}\text{-rating}(eval^{\sigma}, c)$.

In order to determine the dominant strategies, we need to establish how the customers vote, how they are connected, and what the restaurant owner knows. In this paper we assume that the restaurant knows *eval*, leaving the interesting case when *eval* is unknown to future work. We focus instead on the following cases: the restaurant knows the network, the restaurant knows the shape of the network but not the individuals' position, and the network is unknown. We analyse the effect of bribing strategies on $\mathbb{P}$-rating in each such case. Notice how for the case of $\mathbb{O}$-rating the cases collapse to the first. We also look at the special situation in which every customer is a voter.

Given a set of such assumptions, we say that $\mathbb{O}$-rating (or $\mathbb{P}$-rating) are *bribery-proof* under those assumptions if σ^0 is weakly dominant.

2.4 Discussion

Our model is built upon a number of simplifying assumptions which do not play a significant role in the results and could therefore be dispensed with: (i) customers' ratings correspond to their propensity to go to the restaurant. (ii) the restaurant utility equals the sum of all such propensities (iii) bribe $\sigma(c)$ affects evaluation $eval(c)$ linearly. All these assumptions could be generalised by multiplicative factors, such as an average price R paid at the restaurant, and a "customer price" D_c, such that $eval^{\sigma}(c) = eval(c) + \frac{\sigma(c)}{D_c}$.

3 Bribes under $\mathbb{O}$-rating

In this section we look at bribing strategies under $\mathbb{O}$-rating, first focussing on the case where everyone expresses an opinion, then moving on to the more general case.

3.1 All vote

Let us now consider the case in which $V = C$. Recall that $B(\sigma)$ is the set of customers bribed by σ. We say that two strategies σ_1 and σ_2 are *disjoint* if $B(\sigma_1) \cap B(\sigma_2) = \emptyset$. By direct calculation it follows that the revenue of disjoint strategies exhibits the following property:

Lemma 1. *If $V = C$ and σ_1 and σ_2 are two disjoint strategies, then $\mathbf{r}_{\mathbb{O}}(\sigma_1 \circ \sigma_2) = \mathbf{r}_{\mathbb{O}}(\sigma_1) + \mathbf{r}_{\mathbb{O}}(\sigma_2)$.*

We now show that bribing a single individual is not profitable.

Lemma 2. *Let σ be a bribing strategy, $V = C$ and $|B(\sigma)| = 1$. Then, $\mathbf{r}_{\mathbb{O}}(\sigma) \leq 0$, i.e., σ is not profitable.*

Proof sketch. Let $\bar{c}$ be the only individual such that $\sigma(\bar{c}) \neq 0$. By calculation, $\mathbf{r}(\sigma) = u^{\sigma}_{\mathbb{O}} - u^{0}_{\mathbb{O}} = \mathbb{O}\text{-}\mathtt{rating}^{\sigma} - \mathbb{O}\text{-}\mathtt{rating} - \sum_c \sigma(c) = \min\{1, eval(\bar{c}) + \sigma(\bar{c})\} - eval(\bar{c}) - \sigma(\bar{c}) \leq 0$. □

By combining the two lemmas above we are able to show that no strategy is profitable for bribing the $\mathbb{O}$-$\mathtt{rating}$.

Proposition 3. *If $V = C$, then no strategy is profitable.*

Proof sketch. Any bribing strategy σ can be decomposed into n pairwise disjoint strategies such that $\sigma = \sigma_{c_1} \circ \cdots \circ \sigma_{c_n}$ and $|B(\sigma_{c_j})| = 1$ for all $1 \leq j \leq n$. By applying Lemma 1 and Lemma 2 we then obtain that $\mathbf{r}_{\mathbb{O}}(\sigma) \leq 0$. □

From this it follows that σ^0 is weakly dominant and thus $\mathbb{O}$-$\mathtt{rating}$ is bribery-proof when all customers voted.

3.2 Non-voters

Let us now consider the case of $V \subset C$, i.e., when there is at least one customer who is not a voter. In this case Lemma 1 no longer holds, as shown in the following example.

Example 1. *Let $C = \{A, B, C\}$, and let $eval(A) = 0.5$, $eval(B) = 0.5$, and $eval(C) = *$. The initial resources are $u^0 = \mathbb{O}\text{-}\mathtt{rating} \times 3 = 1.5$. Let now $\sigma_1(A) = 0.5$ and $\sigma_1(B) = \sigma_1(C) = 0$, and let $\sigma_2(C) = 0.5$ and $\sigma_2(A) = \sigma_2(B) = 0$. Now $u^{\sigma_1}_{\mathbb{O}} = 0.75 \times 3 - 0.5 = 1.75$ and $u^{\sigma_2}_{\mathbb{O}} = 0.5 \times 3 - 0.5 = 1$, but $u^{\sigma_1 \circ \sigma_2}_{\mathbb{O}} = 0.\bar{6} \times 3 - 1 = 1$.*

The example (in particular σ_1) also shows that $\mathbb{O}$-$\mathtt{rating}$ in this case is not bribery-proof.

We now turn to characterise the set of undominated bribing strategies. We begin by showing that bribing a non-voter is always dominated. Let first σ be a strategy such that $\sigma(\bar{c}) \neq 0$ for some $\bar{c} \in C \setminus V$ and recall that V^{σ} is the set of voters after execution of σ. Let us define the $\bar{c}$-*greedy restriction* of σ to be any strategy $\sigma^{-\bar{c}}$ such that:

- $V^{\sigma^{-\bar{c}}} = V^{\sigma} \setminus \{\bar{c}\}$, i.e., the greedy restriction eliminates $\bar{c}$ from the set of voters.
- For each $c \in V^{\sigma^{-\bar{c}}}$, $\max(1, eval(c) + \sigma(c)) = \max(1, eval(c) + \sigma^{-\bar{c}}(c))$, i.e., the greedy restriction does not waste further resources.
- If there exists $c \in V^{\sigma^{-\bar{c}}}$ such that $eval(c) + \sigma^{-\bar{c}}(c) < 1$ then $\sum_{c \in C} \sigma^{-\bar{c}}(c) = \sum_{c \in C} \sigma(c)$, i.e., the $\sigma^{-\bar{c}}$ redistributes $\sigma(\bar{c})$ among the remaining voters.

We now show that each strategy bribing a non-voter is strictly dominated by any of its greedy restrictions.

Proposition 4. *Let $V \neq C$, and $\bar{c} \in C \setminus V$. Then each strategy σ with $\sigma(\bar{c}) \neq 0$ is strictly dominated by $\sigma^{-\bar{c}}$.*

Proof. Let σ be a strategy with $\sigma(\bar{c}) \neq 0$ for some non-voter $\bar{c}$, and let $\sigma^{-\bar{c}}$ be one of its greedy restriction defined above.

$$u^{\sigma^{-\bar{c}}}_{\mathbb{O}} - u^{\sigma}_{\mathbb{O}} =$$

$$|C|(\mathbb{O}\text{-}\mathtt{rating}^{\sigma^{-\bar{c}}} - \mathbb{O}\text{-}\mathtt{rating}^{\sigma}) + \sum_{c \in C} \sigma(c) - \sum_{c \in C} \sigma^{-\bar{c}}(c) =$$

$$|C|\left(\frac{\sum_{c \in C} eval^{\sigma^{-\bar{c}}}(c)}{|V|} - \frac{\sum_{c \in C} eval^{\sigma}(c)}{|V \cup \bar{c}|}\right) +$$

$$+\left(\sum_{c \in C} \sigma(c) - \sum_{c \in C} \sigma^{-\bar{c}}(c)\right)$$

Observe first that $\sigma^{-\bar{c}}$ is a redistribution, hence $\sum_c \sigma(c) - \sum_c \sigma^{-\bar{c}}(c) \geq 0$, i.e., the second addendum in the above equation is positive. Consider now the case where there exists $c \in V^\sigma \setminus \bar{c}$ such that $eval(c) + \sigma^{-\bar{c}}(c) < 1$. Then by the definition of $\sigma^{-\bar{c}}$ we have that $\sum_{c\in V^\sigma} eval^\sigma(c) = \sum_{c\in V^{\sigma^{-\bar{c}}}} eval^{\sigma^{-\bar{c}}}(c)$, i.e., the greedy restriction preserves the overall evaluation. By straightforward calculation this entails that $u_{\mathbb{O}}^{\sigma^{-\bar{c}}} - u_{\mathbb{O}}^{\sigma} > 0$. If no such c exists, and therefore $\mathbb{O}\text{-}\mathtt{rating}^{\sigma^{-\bar{c}}} = 1$ we have that either $\mathbb{O}\text{-}\mathtt{rating}^{\sigma} < 1$ or, by the efficiency requirement and the fact that $\sigma(\bar{c}) \neq 0$, we have that $\sum_{c\in C} \sigma(c) > \sum_{c\in C} \sigma^{-\bar{c}}(c)$. In either cases we have that $u_{\mathbb{O}}^{\sigma^{-\bar{c}}} - u_{\mathbb{O}}^{\sigma} > 0$. □

Let an $\mathbb{O}$-*greedy strategy* be any efficient strategy that redistributes all the initial resources $u_{\mathbb{O}}^0$ among voters. Making use of the previous result, we are able to characterise the set of all dominant strategies for $\mathbb{O}$-rating.

Proposition 5. *Let $V \neq C$. A strategy is weakly dominant for* $\mathbb{O}$-rating *if and only if it is an $\mathbb{O}$-greedy strategy.*

Proof sketch. For the right-to-left direction, first observe that all $\mathbb{O}$-greedy strategies are payoff-equivalent, and that a non-efficient strategy is always dominated by its efficient counterpart. By Proposition 4 we know that strategies bribing non-voters are dominated, and by straightforward calculations we obtain that in presence of non-voters it is always profitable to bribe as much as possible. For the left-to-right direction, observe that a non-greedy strategy is either inefficient, or it bribes a non-voter, or does not bribe as much as possible. In either circumstance it is strictly dominated. □

While there may be cases in which the number of weakly dominant strategies under $\mathbb{O}$-rating is exponential, all such strategies are revenue equivalent, and Proposition 5 gives us a polynomial algorithm to find one of them: starting from an evaluation vector *eval*, distribute all available resources $u_{\mathbb{O}}^0$ to the voters, without exceeding the maximal evaluation of 1. By either exhausting the available budget or distributing it all, we are guaranteed the maximum gain by Proposition 5.

4 Bribes under $\mathbb{P}$-rating

In this section we look at bribing strategies under $\mathbb{P}$-rating, against various knowledge conditions on the social network. As for Section 3 we start by looking at the case where everyone votes and later on allowing non-voters. Before doing that, we introduce a graph-theoretic measure of influence and prove a useful lemma.

Definition 5. *The* influence weight *of a customer $c \in C$ in a network E and a set of designed voters V is defined as follows:*

$$w_c^V = \sum_{k\in N(c)} \frac{1}{|N(k) \cap V|}$$

Recall that we assumed that every customer can see a voter, thus w_c^V are well-defined for every c. If $V = C$, i.e., when everybody voted, we let $w_c = w_c^C$. In this case, we obtain $w_c = \sum_{k\in N(c)} \frac{1}{deg(k)}$, where $deg(c) = |N(c)|$ is the *degree* of c in E. When V is defined by a bribing strategy σ, we write $w_c^\sigma = w_c^{V^\sigma}$.

Intuitively, each individual's rating influences the rating of each of its connections, with a factor that is inversely proportional to the number of second-level connections that have expressed an evaluation. We formalise this statement in the following lemma:

Lemma 6. *The utility obtained by playing σ with $\mathbb{P}$-$\mathtt{rating}$ is $u^\sigma_{\mathbb{P}} = \sum_{c \in V^\sigma} w^\sigma_c \times eval^\sigma(c) - \sum_{c \in C} \sigma(c)$.*

Proof. By calculation:

$$u^\sigma_{\mathbb{P}} + \sum_{c \in C} \sigma(c) = \sum_{c \in C} \mathbb{P}\text{-}\mathtt{rating}^\sigma(c) = \sum_{c \in C} \underset{k \in N(c) \cap V^\sigma}{\text{avg}} eval^\sigma(k) =$$

$$= \sum_{c \in C} \Big[\frac{1}{|N(c) \cap V^\sigma|} \sum_{k \in N(c) \cap V^\sigma} eval^\sigma(k) \Big] =$$

$$= \sum_{k \in V^\sigma} \Big[\; eval^\sigma(k) \times \sum_{k' \in N(k)} \frac{1}{|N(k') \cap V^\sigma|} \Big] =$$

$$= \sum_{c \in V^\sigma} w^\sigma_c \times eval^\sigma(c)$$

□

4.1 All vote, known network

We begin by studying the simplest case in which the restaurant knows the evaluation *eval*, the network E as well as the position of each customer on the network. The following corollary is a straightforward consequence of Lemma 6:

Corollary 7. *Let $V = C$ and let σ_1 and σ_2 be two disjoint strategies, then $\mathbf{r}_{\mathbb{P}}(\sigma_1 \circ \sigma_2) = \mathbf{r}_{\mathbb{P}}(\sigma_1) + \mathbf{r}_{\mathbb{P}}(\sigma_2)$.*

We are now able to show a precise characterisation of the revenue obtained by any efficient strategy σ (the proof of this proposition can be found in appendix):

Proposition 8. *Let $V = C$, let E be a known network, and let σ be an efficient strategy. Then $\mathbf{r}_{\mathbb{P}}(\sigma) = \sum_{c \in C}(w_c - 1)\sigma(c)$.*

Proposition 8 tells us that the factors w_c are crucial in determining the revenue of a given bribing strategy. Bribing a customer c is profitable whenever $w_c{>}1$ (provided its evaluation was not 1 already), while bribing a customer c with $w_c{\leq}1$ is at most as profitable as doing nothing, as can be seen in the example below. Most importantly, it shows that $\mathbb{P}$-$\mathtt{rating}$ is *not* bribery-proof when the restaurant knows both the network and the customers' evaluations.

Example 2. *Let E be a four-armed star, and let A be the individual in the centre. Assume each individual values the restaurant 0.5. We have that $w_A = 2.2$ and $w_c = 0.7$ for all c different from A. Consider now two bribing strategies: σ^A which bribes A with 0.5, and σ^B which bribes a single individual $B \neq A$ with the same amount. What we obtain is that $\mathbf{r}_{\mathbb{P}}(\sigma^A) = 0.6$, while $\mathbf{r}_{\mathbb{P}}(\sigma^B) = -0.15$.*

Given a network E and an evaluation vector *eval*, let Algorithm 1 define the $\mathbb{P}$-*greedy bribing strategy.*

As a consequence of Proposition 8 we obtain:

Corollary 9. *The $\mathbb{P}$-greedy bribing strategy defined in Algorithm 1 is weakly dominant.*

As in the case of $\mathbb{O}$-$\mathtt{rating}$, Corollary 9 has repercussions on the computational complexity of bribery: it shows that computing a weakly dominant strategy can be done in polynomial time. Notice how the most costly operation lies in the computation of the influence weights w_c, which can be performed only once, assuming the network is static. Similar problems, such as recognising whether bribing a certain individual is profitable, or estimating whether individuals on a network can be bribed above a certain threshold, are also computable in polynomial time.

Input: Evaluation function *eval* and network E
Output: A bribing strategy $\sigma^G_{\mathbb{P}} : C \to Val$

$Budget{=}u^0_{\mathbb{P}}$
$\sigma^G_{\mathbb{P}}(c) = 0$ for all $c \in C$
Compute w_c for all $c \in C$
Sort $c \in C$ in descending order $c_0, \ldots, c_m$ based on w_c

```
for i=0,...,m do
    if Budget≠ 0 then
        if w_{c_i} > 1 then
            σ^G_P(c_i) = min{1 − eval(c_i), Budget}
            Budget=Budget-σ^G_P(c_i)
        end
    end
    return σ^G_P
end
```

Algorithm 1: The $\mathbb{P}$-greedy bribing strategy $\sigma^G_{\mathbb{P}}$

4.2 All vote, unknown network

We now move to study the more complex case of an unknown network. Surprisingly, we are able to show that no bribing strategy is profitable (in expectation), and hence $\mathbb{P}$-rating is bribery-proof in this case. Recall that we are still assuming that the restaurant knows *eval* and everybody voted.

We begin by assuming that the restaurant knows the structure of the network, but not the position of each participant. Formally, the restaurant knows E, but considers any permutation of the customers in C over E as possible. Let us thus define the expected revenue of a strategy σ over a given network E as the average over all possible permutations of customers: $\mathbb{E}[\mathbf{r}_{\mathbb{P}}(\sigma)] = \sum \frac{1}{n!}[u^{\sigma}_{\rho} - u^0_{\rho}]$, where we abuse notation by writing u^{σ}_{ρ} as $u^{\sigma}_{\mathbb{P}}$ under permutation ρ over the network E. What we are able to show is that all strategies are at most as profitable as σ^0 in expected return:

Proposition 10. *Let $V = C$, let the network structure of E be known but not the relative positions of customers on E. Then $\mathbb{E}[\mathbf{r}_{\mathbb{P}}(\sigma)] = 0$ for all strategies σ.*

Proof sketch. Let $|C| = n$. We show the result for any strategy σ that bribes a single customer $\bar{c}$. The general statement follows from the linearity of $\mathbb{E}[\mathbf{r}(\sigma)]$. Equation (5) uses Proposition 8 to compute the revenue for each permutation ρ of customers C on the network:

$$\mathbb{E}[\sigma] = \sum_{\rho} \frac{1}{n!} \ (u^{\sigma}_{\rho} - u^0_{\rho}) = \sum_{\rho} \frac{1}{n!}(w_{\rho(\bar{c})} - 1)\sigma(\bar{c}) = \tag{1}$$

$$= \sum_{c \in C} \frac{(n-1)!}{n!}(w_c - 1)\sigma(\bar{c}) = \frac{(n-1)!}{n!} \sum_{c \in C}(w_c - 1) = 0 \tag{2}$$

The last line follows from the observation that $\sum_c w_c = |C|$ and hence $\sum_c (w_c - 1) = 0$, by a consequence of Definition 5 when everybody votes. □

Hence, if we assume a uniform probability over all permutations of customers on the network, a straightforward consequence of Proposition 10 concludes that it is not profitable (in expectation) to bribe customers.

Corollary 11. *If $V = C$ and the network is unknown, then no bribing strategy for* $\mathbb{P}$-rating *is profitable in expected return.*

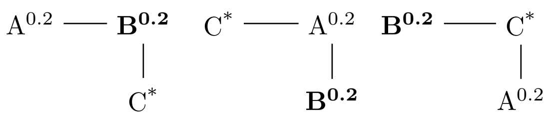

Figure 1: Customers permutations in Example 4.

4.3 Non-voters, known network

With $\mathbb{P}$-rating it is possible to find a network where bribing a non-voter is profitable:

Example 3. *Consider 4 individuals* $\{B, C, D, E\}$ *connected only to a non-voter in the middle. Let* $eval(j) = 0.2$ *for all* j *but the center. We have* $u^0_{\mathbb{P}} = 1$*. Let* A *be the non-voter, and let* $\sigma_1(A) = 1$ *and 0 otherwise. The utility of* σ_1 *is:*

$$\mathbb{P}\text{-rating}^{\sigma_1}(A) + 4\mathbb{P}\text{-rating}^{\sigma_1}(j) - 1 = 1.76$$

All other strategies can be shown to be dominated by σ_1*. Take for instance a strategy* σ_2 *such that* $\sigma_2(B) = 0.8$, $\sigma_2(C) = 0.2$ *and 0 otherwise. The utility of* σ_2 *is* $u^{\sigma_2}_{\mathbb{P}} = 1.25$*.*

It is quite hard to obtain analytical results for strategies bribing non-voters, due to the non-linearity of the $\mathbb{P}$-rating in this setting. We can however provide results in line with those of the previous section if we restrict to *voter-only strategies*, i.e., strategies σ such that $\sigma(c) = 0$ for all $c \notin V$. In this case, a similar proof to Proposition 8 shows the following:

Proposition 12. *Let* $V \neq C$, E *be a known network, and* σ *be an efficient bribing strategy such that* $B(\sigma) \subseteq V$*. Then,* $\mathbf{r}_{\mathbb{P}}(\sigma) = \sum_{c \in V}(w^V_c - 1)\sigma(c)$*.*

The difference with the case of $V = C$ is that w^V_c can be arbitrarily large in the presence of non-voters, such as in our Example 3.

4.4 Non-voters, unknown positions

Unlike the case of $V = C$, in this case it is possible to define bribing strategies that are profitable (in expected return).

Example 4. *Let* $C = \{A, B, C\}$*, and the initial evaluation* $eval(A) = eval(B) = 0.2$ *and* $eval(C) = *$*. Assume that the structure of the network is known, but the position of the individuals is not. Let the three possible network positions (without counting the symmetries) be depicted in Figure 1. Let* $\sigma(B) = 0.2$ *and* $\sigma(A) = \sigma(C) = 0$*. In the first case:*

$$\begin{aligned}\mathbf{r}^1_P(\sigma) =& \mathbb{P}\text{-rating}(A) + ... + \mathbb{P}\text{-rating}(C) - 0.2 - u^0_{\mathbb{P}} = \\ &= 0.3 + 0.3 + 0.4 - 0.2 - 0.6 = 0.2\end{aligned}$$

In the second case $\mathbf{r}^2_{\mathbb{P}}(\sigma) = 0$ *while in the third:*

$$\mathbf{r}^3_P(\sigma) = 0.4 + 0.3 + 0.2 - 0.2 - 0.6 = 0.1$$

Therefore, $\mathbb{P}$-rating is not bribery-proof (in expectation) in the presence of non-voters when the network is unknown. Interesting computational problems open up in this setting, such as identifying the networks that allow for profitable bribing strategies, and their expected revenue.

5 Boundaries of bribery-proofness

The previous sections have shown that having a network-based rating systems, where individuals are influenced by their peers, is not bribery-proof, even when the position of individuals in a given network is not known. However bribing strategies have a different effect in the overall score. While the utility of $\mathbb{O}$-$\mathtt{rating}$ is a sum of the *global* average of voters' evaluation, the utility of $\mathbb{P}$-$\mathtt{rating}$ is a sum of *local* averages of voters' evalution against the one of their peers.

Therefore a strategy bribing one voter affects everyone in the case of $\mathbb{O}$-$\mathtt{rating}$, but it can be shown to have a limited effect in the case of $\mathbb{P}$-$\mathtt{rating}$.

Proposition 13. *Let σ be an efficient strategy s.t. $|B(\sigma)| = 1$, and let $\bar{c}$ be such that $\sigma(c) \neq 0$. Then $\mathbf{r}_{\mathbb{P}}(\sigma) < N(\bar{c})$.*

Proof. By calculation, we have that:

$$\mathbf{r}_{\mathbb{P}}(\sigma) = \sum_{c \in C} \mathbb{P}\text{-}\mathtt{rating}^{\sigma}(c) - \sigma(\bar{c}) - \sum_{c \in C} \mathbb{P}\text{-}\mathtt{rating}(c) =$$

$$\sum_{c' \in N(\bar{c})} \mathbb{P}\text{-}\mathtt{rating}^{\sigma}(c) - \sigma(\bar{c}) - \sum_{c' \in N(\bar{c})} \mathbb{P}\text{-}\mathtt{rating}(c) \leq$$

$$\leq 1 \times N(\bar{c}) - \sigma(\bar{c}) - \sum_{c' \in N(\bar{c})} \mathbb{P}\text{-}\mathtt{rating}(c) < N(\bar{c})$$

□

The previous result shows that increasing the number of individuals that are not connected to an agent that is bribed, even if these are non-voters, does not increase the revenue of the bribing strategy. This is not true when we use $\mathbb{O}$-$\mathtt{rating}$.

Proposition 14. *Let σ be an efficient strategy. The revenue $\mathbf{r}_{\mathbb{O}}(\sigma)$ of σ is monotonically increasing with the number of non-voters, and is unbounded.*

Proof. It follows from our definitions that:

$$\mathbf{r}_{\mathbb{O}}(\sigma) = (\frac{|C|}{|V^{\sigma}|} - 1)\big[\sum_{c \in C} eval(c) + \sigma(c)\big]$$

The above figure is unbounded and monotonically increasing in the number of non-voters, which can be obtained by increasing C keeping V^{σ} fixed. □

So while $\mathbb{P}$-$\mathtt{rating}$ and $\mathbb{O}$-$\mathtt{rating}$ are not bribery-proof in general, it turns out that the impact of the two in the overall network are significantly different. In particular, under realistic assumptions such as a very large proportion of non-voters and with participants having a few connections, bribing under $\mathbb{O}$-$\mathtt{rating}$ is increasingly rewarding, while under $\mathbb{P}$-$\mathtt{rating}$ this is no longer the case.

6 The Complexity of Bribery under the $\mathbb{P}$-$\mathtt{rating}$

We now investigate, from a complexity theoretic standpoint, the problem of computing a bribing strategy yielding at least some given revenue, when not every customer votes and the restaurant has full knowledge of each customer's position. Firstly we re-formulate the above optimisation problem as a decision problem.

BRIBE-NVKL
Instance: Network (C, E), evaluation $eval_0$, $\rho \in \mathbb{Q}$
Yes-Instance: An instance of BRIBE-NVKL s.t. there exists a strategy σ with $\mathbf{r}(\sigma) \geq \rho$

Any instance of the above problem should adhere to the usual restrictions of the framework. These are, most importantly, that the initial evaluation is such that every customer $c \in C$ is adjacent to at least one customer $c' \in C$ such that $eval(c') \neq *$ (recall that every customer is adjacent to itself). Also, any strategy σ is such that $\sum_{c \in C} \sigma(c)$ is at most the initial utility resulting from $eval_0$.

The following proposition is straightforward:

Proposition 15. BRIBE-NVKL *is in* **NP**.

Proof. Given a customer network (C, E), an evaluation *eval*, and $\rho \in \mathbb{Q}$, we can clearly decide whether a given strategy σ yields a revenue of at least ρ in polynomial-time (we simply evaluate the strategy). It therefore follows that BRIBE-NVKL is in **NP**. □

In what follows, we show that BRIBE-NVKL is **NP**-hard, by giving a reduction from the known **NP**-complete problem of finding an independent set on 3-regular graphs, aka ISREG(3) Garey and Johnson (1990).

Recall that a graph G is 3-regular if the degree of every vertex is 3, and an independent set of G is a subset X of its vertices such that there is no edge of G joining any pair of vertices in X. We can now give the following definition:

ISREG(3)
Instance: A 3-regular graph G, $k \in \mathbb{N}$
Yes-Instance: An instance of ISREG(3) such that G has an independent set of size at least k

The proof of the following proposition can be found in appendix:

Proposition 16. BRIBE-NVKL *is* **NP**-*hard.*

As a direct consequence of Propositions 15 and 16 we obtain the following:

Theorem 17. BRIBE-NVKL *is* **NP**-*complete.*

In summary, we have been able to prove the **NP**-completeness of BRIBE-NVKL by giving a reduction from ISREG(3). This is an important finding, that significantly strengthens the value of personalised rating systems and their resistance to bribery, as we have demonstrated that we cannot compute an optimal bribing strategy, nor any strategy guaranteeing at least a given reward, in a reasonable amount of time; that is, of course, unless **P** = **NP**.

7 Conclusive remarks

We introduced $\mathbb{P}$-rating, a network-based rating system which generalises the commonly used $\mathbb{O}$-rating, and analysed their resistance to external bribery under various conditions. The main take-home message of our contribution can be summarised in one point, deriving from our main results:

$\mathbb{P}$-rating and $\mathbb{O}$-rating are not bribery-proof in general. However, if we assume that a service provider has a cost for bribing an individual, there are situations in which $\mathbb{P}$-rating is fully bribery proof, while $\mathbb{O}$-rating is not. For instance, if the cost of bribing an individual c is at least $N(c)$ then $\mathbb{P}$-rating is bribery-proof. As observed previously, this is not necessarily true for $\mathbb{O}$-rating. In particular, if we assume the presence of unreachable individuals the difference is more significant. As shown, for $\mathbb{P}$-rating we need to bribe individuals with $w_c > 1$. With $\mathbb{O}$-rating is sufficient to find one voter who accepts a bribe.

To strengthen for the practical applicability of the personalised rating framework, we also showed that despite the proposed system being manipulable such task is intractable in practice, as the problem of computing the existence of a manipulation strategy guaranteeing a given reward and thus an optimal one, what we called BRIBE-NVKL, is **NP**-complete.

References

N. Alon, M. Feldman, O. Lev, and M. Tennenholtz. How robust is the wisdom of the crowds? In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI)*, 2015.

R. Andersen, C. Borgs, J. Chayes, U. Feige, A. Flaxman, A. Kalai, V. Mirrokni, and M. Tennenholtz. Trust-based recommendation systems: An axiomatic approach. In *Proceedings of the 17th International Conference on World Wide Web (WWW)*, 2008.

K. R. Apt and E. Markakis. Social networks with competing products. *Fundamenta Informaticae*, 129(3):225–250, 2014.

D. Baumeister, G. Erdélyi, and J. Rothe. How hard is it to bribe the judges? A study of the complexity of bribery in judgment aggregation. In *Proceedings of the Second International Conference on Algorithmic Decision Theory (ADT)*, 2011.

F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2015.

R. Bredereck, J. Chen, S. Hartung, S. Kratsch, R. Niedermeier, O. Suchý, and G. J. Woeginger. A multivariate complexity analysis of lobbying in multiple referenda. *Journal of Artificial Intelligence Research*, 50:409–446, 2014.

R. Bredereck, P. Faliszewski, R. Niedermeier, and N. Talmon. Large-scale election campaigns: Combinatorial shift bribery. *Journal of Artificial Intelligence Research (JAIR)*, 55:603–652, 2016.

M. Brill, V. Conitzer, R. Freeman, and N. Shah. False-name-proof recommendations in social networks. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2016.

N. Bu. Unfolding the mystery of false-name-proofness. *Economics Letters*, 120(3):559–561, 2013. ISSN 0165-1765.

R. Christian, M. Fellows, F. Rosamond, and A. Slinko. On complexity of lobbying in multiple referenda. *Review of Economic Design*, 11(3):217–224, 2007.

V. Conitzer, N. Immorlica, J. Letchford, K. Munagala, and L. Wagman. False-Name-Proofness in Social Networks. In *6th International Workshop on Internet and Network Economics (WINE)*, 2010.

R. Conte and M. Paolucci. *Reputation in Artificial Societies: Social Beliefs for Social Order*. Kluwer Academic Publishers, 2002.

R. Conte, M. Paolucci, and J. Sabater-Mir. Reputation for innovating social networks. *Advances in Complex Systems*, 11(2):303–320, 2008.

P. Faliszewski and R. Niedermeier. Parameterization in computational social choice. In M.-Y. Kao, editor, *Encyclopedia of Algorithms*. Springer Berlin Heidelberg, 2014.

P. Faliszewski, E. Hemaspaandra, and L. A. Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research (JAIR)*, 35:485–532, 2009.

F. Garcin, B. Faltings, and R. Jurca. Aggregating reputation feedback. In *Proceedings of the International Conference on Reputation*, 2009.

M. R. Garey and D. S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1990.

U. Grandi. Social choice and social networks. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 9, pages 169–184. AI Access, 2017.

U. Grandi and P. Turrini. A network-based rating system and its resistance to bribery. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-2016)*, July 2016.

U. Grandi, J. Stewart, and P. Turrini. The complexity of bribery in network-based rating systems. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2018.

E. Helpman and T. Persson. Lobbying and legislative bargaining. Working Paper 6589, National Bureau of Economic Research, June 1998.

O. Lev and M. Tennenholtz. Group recommendations: Axioms, impossibilities, and random walks. *CoRR*, abs/1707.08755, 2017. URL `http://arxiv.org/abs/1707.08755`.

I. Pinyol and J. Sabater-Mir. Computational trust and reputation models for open multi-agent systems: a review. *Artificial Intelligence Review*, 40(1):1–25, 2013.

J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1):33–60, 2005.

S. Simon and K. R. Apt. Social network games. *Journal of Logic and Computation*, 25(1): 207–242, 2015.

T. Todo and V. Conitzer. False-name-proof matching. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2013.

Umberto Grandi
IRIT, University of Toulouse
Toulouse, France
Email: `umberto.grandi@irit.fr`

James Stewart
Amadeus IT Group
Nice, France
Email: `james.stewart13@imperial.ac.uk`

Paolo Turrini
University of Warwick
Warwick, U.K.
Email: `p.turrini@warwick.ac.uk`

Appendix

Proof of Proposition 8:

Proof. By calculation, where Step (2) uses Lemma 6, and Step (4) uses the fact that σ is efficient:

$$\mathbf{r}_{\mathbb{P}}(\sigma) = u^{\sigma}_{\mathbb{P}} - u^{0}_{\mathbb{P}} = \tag{3}$$

$$= [\sum_{c\in C} w_c \; eval^{\sigma}_c - \sum_{c\in C} \sigma(c) - \sum_{c\in C} w_c \; eval(c)] = \tag{4}$$

$$= \sum_{c\in C} \left[w_c \left[\min\{1, eval(c) + \sigma(c)\} - eval(c) \right] \right] - \sum_{c\in C} \sigma(c) \tag{5}$$

$$= \sum_{c\in C} (w_c - 1)\sigma(c). \tag{6}$$

□

Proof of Proposition 16:

Proof. We start by giving a reduction from an arbitrary instance of ISREG(3) to an instance of BRIBE-NVKL. That is, given a 3-regular graph G and $k \in \mathbb{N}$, we construct a network (C, E), an initial evaluation $eval_0$, and $\rho \in \mathbb{Q}$ such that G has an independent set of size at least $k \iff$ there exists a strategy on $((C, E), eval_0)$ that yields a revenue of at least ρ. Given a 3-regular graph G, we define a network of customers as follows:

Customers The set C of customers is composed of *old*, *pendant*, and *edge* cutomers. For all vertices $v \in G$, we create an *old customer* $v \in C$, as well as a set of *pendant customers* $v_1, .., v_n \in C$, where n is the number of vertices of G. For each edge (u, v) of G, we introduce an *edge customer* $w_{u,v} \in C$.

Network The network E relating customers is defined as follows. For each old customer v, there is an edge $(v, v_i) \in E$ for $i = 1, 2, ..., n$, connecting it to the related pendant customers. For every edge (u, v) of G, we add $(u, w_{u,v})$ and $(w_{u,v}, v)$ to E, relating the two old customers with the corresponding edge customer.

For any such network as constructed above, we can define an initial evaluation $eval_0$ as follows, where $0 < \epsilon < 1$ is some value that will be set later in the proof:

- If $c \in C$ is an old customer then $eval_0(c)$=$*$ (non-voter).
- If $c \in C$ is an edge or pendant customer then $eval_0(c)$=ϵ.

An example of the construction of the customer network and evaluation from a graph G can be seen in Figure 2.

By the construction of the network, we have that for all $c \in C$, the $\mathbb{P}$-rating$(c, eval_0) = \epsilon$ (recall that we assumed $c \in N(c)$ for all customers). Every customer of the newly constructed customer network contributes ϵ to the initial utility of the network and therefore $u^0_{\mathbb{P}} = \epsilon(n + n^2 + \frac{3n}{2})$. We now choose ϵ so that $u^0_{\mathbb{P}} = k$; that is, so that

$$\epsilon = \frac{k}{n + n^2 + \frac{3n}{2}}.$$

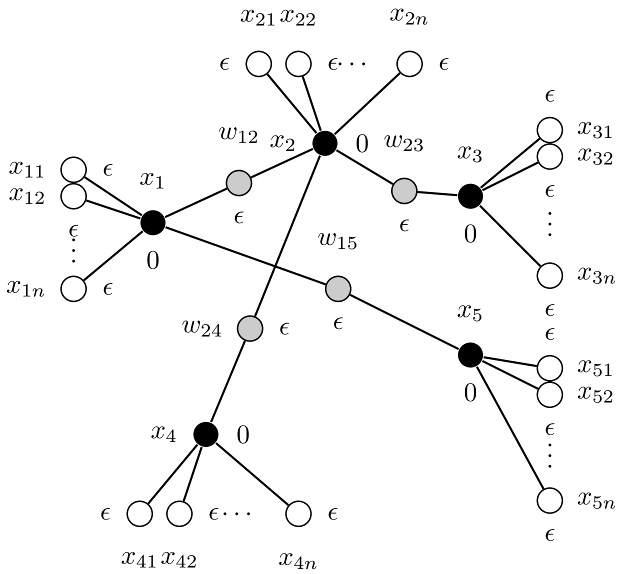

Figure 2: The figure above shows a portion of a 3-regular graphs, formed by the five black vertices (additional edges required by 3-regularity have been omitted). Black vertices are therefore old customers, white vertices are pendant customers, and grey vertices are edge customers. The associated bribing strategy is marked with 0, 1 and ϵ labels.

By assumption the restaurant owner can only make bribes totalling at most k. Furthermore, note that the initial evaluation is a valid one in that every customer of the network is adjacent to at least one voter. Finally, let

$$\rho = k(1-\epsilon)\left(\frac{1}{n+4} + \frac{n+3}{2}\right) - k.$$

($\implies$) Suppose that our instance (G, k) of ISREG(3) is a yes-instance; that is, there is a set I of k vertices such that no two vertices of I are adjacent in G. Consider the bribing strategy for (C, E) (as constructed above) where $\sigma(c) = 1$, for every old customer corresponding to some vertex of I, and $\sigma(c') = 0$ for all other $c' \in C$.

Let us now compute the revenue obtained by σ. Recall that the revenue is equal to the increase in $\mathbb{P}$-rating of the bribed customers and their neighbourhoods (old, pendant, and edge customers), minus the cost of the bribe. The cumulative increase in rating of bribed old customers is:

$$k\left(\frac{1+(n+3)\epsilon}{n+4} - \epsilon\right) = k\frac{1-\epsilon}{n+4}.$$

The cumulative increase in rating of pendent customers is:

$$nk\left(\frac{1+\epsilon}{2} - \epsilon\right) = nk\frac{1-\epsilon}{2}.$$

Finally, the increase in rating due to edge customers is:

$$3k\left(\frac{1+\epsilon}{2} - \epsilon\right) = 3k\frac{1-\epsilon}{2}.$$

Recall that bribed old customers correspond to an independent set in G. Summing up, the revenue of strategy σ is:

$$k(1-\epsilon)\left(\frac{1}{n+4}+\frac{n+3}{2}\right)-k=\rho.$$

Therefore $((C,E), eval_0, \rho)$ is a yes-instance of NVKL.

($\Longleftarrow$) We now suppose that $((C,E), eval_0, \rho)$ is a yes-instance of NVKL and that σ is a bribing strategy that yields a revenue of at least ρ. We will assume that σ is also optimal, i.e., that there is no other strategy σ' yielding a higher revenue. We will now show that σ can be transformed into a revenue-equivalent strategy such that (a) only old customers are bribed, (b) all bribed old customers are bribed fully, and (c) exactly k old customers are bribed.

We begin by showing the following technical lemma, whose proof is omitted in the interest of space:

Lemma 18. *Let σ be an optimal bribing strategy. Let X be the set of customers for which $eval(x) < eval^{\sigma}(x) < 1$. Let $v_y = |N(y) \cap V|$ for any $y \in C$. For $x, y \in X$*

$$\sum_{z \in N(x)} \frac{1}{v_z} = \sum_{z \in N(y)} \frac{1}{v_z}.$$

So, given an optimal bribing strategy, we can move bribes amongst non-fully bribed voters arbitrarily without affecting the revenue acquired so long as we do not totally remove all the bribe from a customer that was not originally a non-voter, and we do not turn a non-voter into a voting one.

Revenue equivalent strategy - new customers. Let us call *new customers*, the set of edge and pendant customers. We begin by showing that σ can be modified into an optimal strategy that does not bribe any new customer.

If a new customer is bribed then the bribes to new customers can be enumerated in descending order as $1-\epsilon, 1-\epsilon, ..., 1-\epsilon, \epsilon_1, \epsilon_2, ..., \epsilon_s$, for some $s \geq 0$ and where $0 < \epsilon_i < 1-\epsilon$ for each $i = 1, 2, ..., s$, with possibly no bribe of $1-\epsilon$. By Lemma 18, we can move bribes amongst the new customers so that we may assume that all but at most one new customer is not fully bribed; that is, that $s \leq 1$. The following result is needed:

Lemma 19. *There exists an old customer c' who has not been bribed and where at most one of its adjacent new pendant customers has been bribed.*

First, we suppose that *there exists a fully bribed new customer*, and derive a contradiction with the optimality of σ.

Consider the bribe of $1-\epsilon$ to c, and consider the increase in $\mathbb{P}$-rating generated by this single bribe. If c is a new pendant customer then this contribution is certainly less than 2 as $|N(c)| = 2$, and if c is a new edge customer then this contribution is less than 3 as $|N(c)| = 3$. Therefore, in all cases, the bribe of $1-\epsilon$ to c contributes less than 3 units to the overall utility accrued from σ.

By Lemma 19, let c' be an old customer that is not bribed and that is adjacent to at most one new pendant customer that has been bribed. Consider moving the $1-\epsilon$ bribe from c to c'; so, we obtain a new (efficient) strategy σ'. Let us examine the increase in $\mathbb{P}$-rating generated by this new $1-\epsilon$ bribe.

At least $n-1$ of the new pendant customers adjacent to c' have not been bribed and so the associated cumulative increase in rating is given by $(n-1)\frac{1}{2}-(n-1)\epsilon$ and given that

$\epsilon \leq \frac{2}{2n+5}$ then the cumulative increase in utility is

$$(n-1)\left(\frac{1}{2}-\epsilon\right) > \frac{n-1}{2} - 1.$$

Bribing c' might reduce the $\mathbb{P}$-ratings of c' and its adjacent new edge customers. However, this reduction is certainly less than 4 units. Therefore we may conclude that the movement of $1-\epsilon$ of bribe from c to c' increases the overall utility by an amount greater than $\left(\frac{n-1}{2}-1\right)-7$ units. This amount is strictly positive for n sufficiently large ($n \geq 14$). Therefore the strategy σ' that we have constructed yields a revenue greater than that of σ, in contradiction with its optimality.

Suppose now that *some new customer c has been bribed some amount δ such that* $0 < \delta < 1-\epsilon$. By a detailed case study – omitted for space constraints – we can again derive a contradiction with the optimality of σ. Therefore, we conclude that *no new customer have been bribed* in the revenue-equivalent optimal strategy σ.

Revenue equivalent strategy - old customers. We now turn our attention to old customers. The bribes on old customers can be enumerated in descending order as $1, 1, ..., 1, \delta_1, \delta_2, ..., \delta_m$, for some $m \geq 0$ and where $0 < \delta_i < 1$, for each $i = 1, 2, ..., m$, with possibly no bribes of 1. Without loss of generality, we may assume that $\sum_{i=1}^{m} \delta_i \leq 1$; otherwise, we would have that $m \geq 2$ and we could reduce the bribes $\delta_2, \delta_3, ..., \delta_m$, without making any equal to zero, so as to increase the bribe δ_1 to 1 and secure another fully bribed customer. The following result is needed:

Lemma 20. *Let (C, E) be some network with initial evaluation $eval_0$ and let σ be a bribing strategy. Let $c \in C$ be such that $eval_0(c) \neq *$ and $eval^{\sigma}(c) = \delta > 0$, but where for every customer $c'' \in \bigcup\{N(c') : c' \in N(c)\}$, we have that $\delta < eval^{\sigma}(c'')$. If σ_{-c} is the bribing strategy obtained from σ by removing the bribe from c, we have that $\mathbf{r}(\sigma_{-c}) \geq \mathbf{r}(\sigma)$.*

Therefore, we can assume that at most one old customer has not been fully bribed.

Suppose now that there is in fact *one bribed old customer that has not been fully bribed.* Let us call this old customer c and further suppose that it has been bribed δ where $0 < \delta < 1$. We will again show that this yields yet another contradiction with the optimality of σ. We have the capacity to increase this bribe to 1 at a cost of $1-\delta$ (which we can do, given the remaining resource). The $\mathbb{P}$-rating of all the customers within $N(c)$ will increase with the cumulative increase (only due to new pendant neighbours) being

$$n\frac{1+\epsilon}{2} - n\frac{\delta+\epsilon}{2} = n\frac{1-\delta}{2}.$$

Hence we obtain an increase in revenue for n sufficiently large ($n \geq 3$). This contradicts the optimality of σ. Henceforth, we assume that, without loss of generality, any optimal bribing strategy σ on (C, E), with initial evaluation $eval_0$, is necessarily such that only old customers are bribed and bribed old customers are fully bribed.

Suppose now that the bribing strategy σ *bribes less than k old customers*; so, there is an old customer c that has not been bribed. Let us amend σ to obtain a new bribing strategy σ' by bribing c so that $\sigma(c) = 1$. This costs us 1 unit of resource. There is no customer of C such that its $\mathbb{P}$-rating decreases, and the cumulative increase in $\mathbb{P}$-rating of the n new pendant customers adjacent to c is

$$n\left(\frac{1+\epsilon}{2}-\epsilon\right) = n\left(\frac{1-\epsilon}{2}\right) > \frac{n(2n+3)}{2(2n+5)} > \frac{n}{4}$$

which is strictly greater than 1 (the amount invested) for n sufficiently large ($n \geq 5$). This contradicts the optimality of σ. Furthermore, it is clear that more than k old customers could not have been bribed since the initial utility of the network totals only k and each old customer is bribed by 1.

Finding an independent set of size k. We have shown above that the optimal bribing strategy σ on (C, E) is such that only old customers are bribed, all bribed old customers are fully bribed, and exactly k old customers are bribed.

Consider now the revenue accruing from our optimal bribing strategy σ. Irrespective of which k old customers are fully-bribed, the increase in $\mathbb{P}$-rating due to these old customers is equal to:

$$\frac{(1+(n+3)\epsilon)}{n+4} - \epsilon = \frac{1-\epsilon}{n+4},$$

and the $\mathbb{P}$-rating of the pendant customers adjacent to each of these bribed old customers increases by:

$$\frac{1+\epsilon}{2} - \epsilon = \frac{1-\epsilon}{2}.$$

All that remains is to compute the revenue accruing due to the new edge customers adjacent to each of these bribed old customers (as the $\mathbb{P}$-rating of any other old or new customer does not change). However, this depends upon how many bribed old customers each new edge customer is adjacent to. Let m_i denote the number of new edge customers adjacent to i bribed old customers, for $i = 1, 2$. If a new edge customer c is adjacent to 1 bribed old customer then its increase in $\mathbb{P}$-rating is

$$\frac{(1+\epsilon)}{2} - \epsilon = \frac{(1-\epsilon)}{2}$$

and if it is adjacent to two bribed old customers then its increase in $\mathbb{P}$-rating is

$$\frac{(2+\epsilon)}{3} - \epsilon = \frac{2(1-\epsilon)}{3}.$$

So, the total increase in revenue is

$$m_1 \frac{(1-\epsilon)}{2} + m_2 \frac{2(1-\epsilon)}{3}.$$

We also know that by counting the edges joining bribed old customers and their adjacent new edge customers, we obtain that $3k = 2m_2 + m_1$. Hence, the total increase in $\mathbb{P}$-rating due to new edge customers is equal to

$$\begin{aligned} & m_1 \frac{(1-\epsilon)}{2} + m_2 \frac{2(1-\epsilon)}{3} \\ &= (3k - 2m_2)\frac{(1-\epsilon)}{2} + m_2 \frac{2(1-\epsilon)}{3} \\ &= \frac{3k(1-\epsilon)}{2} - m_2 \frac{(1-\epsilon)}{3}. \end{aligned}$$

So, the revenue due to the bribing strategy σ is:

$$\begin{aligned} & \frac{k(1-\epsilon)}{n+4} + \frac{nk(1-\epsilon)}{2} + \frac{3k(1-\epsilon)}{2} - m_2 \frac{(1-\epsilon)}{3} - k \\ = \quad & (1-\epsilon)[\frac{k}{n+4} + \frac{k(n+3)}{2} - \frac{m_2}{3}] - k. \end{aligned}$$

Clearly this revenue is largest when m_2 is 0, and if $m_2 > 0$ then the revenue is less than this maximal value. Also, when m_2 is 0 this revenue is exactly equal to ρ. Hence, as we started with a yes-instance of NVKL, we must have that $m_2 = 0$, i.e., that no edge customer is adjacent to two bribed old customers. Thus, the k vertices of G corresponding to the k bribed old customers in C form an independent set, and (G, k) is a yes-instance of ISREG(3). $\square$