The Reassignment Problem

Lenore Cowen Tara Kola Eli Rosmarin
Department of Computer Science
Tufts University
Medford, MA 02155
Correspondence to: cowen@cs.tufts.edu

The reassignment problem is our name for a set of assignment problems that are defined in reference to a pre-existing initial assignment. For example, for a set of n people assigned to n houses, we consider reassignments that either never place a person into a house they like less than their currently assigned house, or alternatively consider reassignments where a bound is placed on the number of people whose new assignment violates this constraint. We find that budgeted weighted perfect re-matching is NP-hard, but we are able to give a polynomial time algorithm for budgeted unweighted perfect re-matching. We connect this work to related literature that considers markets for barter and exchange, discuss generalizations, and open problems.

1 INTRODUCTION

In the classical assignment problem, there are n people and n houses. We place an edge between person p and house h if p can be assigned to h. Then there is an assignment if and only if this bipartite graph has a perfect matching [Hall, 1935]. The classical assignment problem is typically generalized in two ways: in the weighted perfect assignment problem [Kuhn, 1955, Munkres, 1957], edges have weights, and the goal is to generate a perfect matching of maximum weight. In the budgeted (weighted) assignment problem, edges also have costs, and the goal is to determine if a perfect matching exists whose total cost is within the budget (or to find such an assignment of maximum weight) [Berger et al., 2008]. Classical and weighted perfect assignment problems are polynomial-time solvable [Lovász and Plummer, 2009], but the weighted budgeted versions are NP-hard [Berger et al., 2008]. The subject of this paper is a particular set of constrained assignment problems that we call the *reassignment problem.* In the reassignment problem, we assume that p_i is already occupying some house h_i, and will only release h_i back into the set of houses that can be assigned to others if they are matched to a house that they prefer to h_i or, in some variants of the problem, if they are offered alternative compensation for accepting assignment to a house they prefer less. We formulate several different unweighted and weighted versions of the re-matching problem in this paper, and consider some budgeted extensions, analogous to the classical assignment problem. Some of these re-matching problems we show can be solved using already known algorithms from matching theory; some of them we prove are computationally difficult; and some whose complexity seems unknown.

Our motivation for the reassignment problem comes from that of a *market maker* who wants to profit from forming an exchange market. This market maker either gets compensated by a fixed fee from each participant in the exchange market, or alternatively, receives a reward proportional to the number of people who are successfully re-matched. People, who already own houses, would be motivated to participate in the exchange market because depending on the structure of the preference graph, they might be assigned a house that they consider better than their current house, and they are guaranteed to never be assigned a worse house (or alternatively, are acceptably compensated when they are assigned a worse house.) While we use the language of people and houses throughout this paper, such a scenario might be applicable in any market with limited liquidity where desired goods are owned and unique or rare, such as fine art objects. For another

Manuscript submitted for review to ACM Economics & Computation 2018 (EC '18).

different application of the reassignment problem: consider a supervisor who is assigning weekly shifts to a workforce. Some people may prefer early morning shifts, some may prefer weekend shifts, but suppose most people don't like late night shifts (but a small percent do). In order to be fair, an initial assignment is made by lottery. However, in order to increase overall employee happiness, the employer is willing to engage in re-matching, so that employees can swap for better shifts as long as everyone agrees to the swaps (or pay additional compensation to employees who cover shifts that no one wants). We now define our problems formally and commence an examination of their computational complexity.

2 PROBLEM DEFINITIONS

Unweighted Perfect Rematching Suppose P is a set of n people and H is a set of n houses, where we re-index the houses so that person p_i is currently matched with house h_i. Associated to each person is a (possibly empty) subset of H, represented the set of houses that they prefer to their current house. Find the assignment f that maximizes the number of re-matches, where $f : P \rightarrow H$ matches people to houses such that 1) f is a bijection, and 2) f assigns each person either to their current house, or a house that they prefer to their current house.

Lemma 2.1. *Unweighted perfect rematching can be solved in polynomial time*

Proof. Follows directly from the fact that there exist polynomial time algorithms for maximum weighted matching in a bipartite graph. Set positive $\epsilon < 1/n$. We form a bipartite graph $G = (P \cup H, E)$, with edges e_{ii} of weight ϵ and for $i \neq j$ we place an edge $e_{ij} \in E$ of weight 1 whenever p_i prefers h_j to h_i. Then a re-matching that gives r people a better assignment corresponds to a perfect matching of weight $r + (n - r)\epsilon < r + 1$, so solutions to the unweighted perfect rematching problem correspond to maximum weight perfect matchings in this graph. □

Scaling all weights in the proof of the Lemma above so that they become integral, we can obtain integral weights of size bounded by $O(n)$, so using either the algorithm of [Gabow and Tarjan, 1991], or [Orlin and Ahuja, 1992] (see also Table II of [Duan and Pettie, 2014]), we can compute unweighted perfect rematching in $O(n^{2.5} \log n)$ time.

The above formulation, doesn't distinguish whether a person is assigned to a house that they consider only slightly better or massively better than the house they were originally assigned; in fact, if they are only assigned to an incrementally better house, they are more likely to participate in the re-matching process again; in the model where the market maker is rewarded based on the *number* of participants whose happiness increases, rather than a percentage of the increase, this model makes sense. On the other hand, if the market maker is instead awarded a *percentage* of the increase, we instead want to allow edge weights, where the weight is proportional to the bonus that the person will pay the market maker if they succeed in making that re-match.

Perfect Rematching. Suppose P is a set of n people and H is a set of n houses, where we re-index the houses so that person p_i is currently matched with house h_i. Associated to each person k is a (possibly empty) subset of H, H_k representing the set of houses that they prefer to their current house, and for each $h_j \in H_k$, let $w_{ij} > 0$ represent the reward that p_i will pay to be matched with house j instead of their current house. Find the assignment f that maximizes the sum of the weights of the re-matches, where $f : P \rightarrow H$ matches people to houses such that 1) f is a bijection, and 2) f assigns each person either to their current house, or a house that they prefer to their current house.

If we assume all the w_{ij} are > 1, or alternatively, choose a new ϵ such that $\epsilon < (\min w_{ij})/n$, then the problem again reduces to a special case of weighted perfect matching in a bipartite graph, the same as the unweighted case. Thus weighted perfect re-matching can also be solved in polynomial time.

So far, the versions of the re-matching problem we have considered can be easily solved by a simple transformation to weighted bipartite matching. Thus the main contribution of this paper concerns considering budgeted versions of these problems. It may be, that the current matching is maximum or near maximum among all perfect matchings; in this case, it is hard for the market maker to create an exchange market that will have any transactions without relaxing the constraint that people can never be matched with houses they prefer less than their current house. Instead, perhaps the market maker can put liquidity into the system as follows: the market maker can compensate people to accept houses that they like *less* than their current house. It will be application dependent whether the compensation consists of additional desirable housing stock that the market maker owns, or monetary compensation that allows the compensated individuals to rent or buy outside of the confines of the exchange market.

For example, in the case of houses, maybe the market maker buys or builds some number of additional vacant houses that are extremely nice and spacious that everyone would prefer to their current house; thus some number of people can be left unmatched with the thought that they will move into these new vacant houses. Or in the case of a fine art auction, maybe the market maker buys some Degas sketches that it can offer into the market, or alternatively, promises people who do not receive an art object they consider better than the one that they are willing to exchange, will each be paid $200,000. All these variations add in some flexibility that may allow us to unlock more value (in some instances substantially more value) from the exchange market, but in each case, there is some bound of *budget* or cost associated with making certain assignments. Thus we define the following budgeted versions of the re-matching problems:

Budgeted Unweighted Perfect Rematching: Version I. Suppose P is a set of n people and H is a set of n houses, where we re-index the houses so that person p_i is currently matched with house h_i, and $B > 0$ is a positive integer. Associated to each person is a (possibly empty) subset of H, representing the set of houses that they prefer to their current house. Find the assignment f that maximizes *the number of people re-matched to a house they prefer to their current house,* where $f : P \to H$ matches people to houses such that 1) f is a bijection, and 2) f assigns at most B people to a house they prefer less than their current house.

Budgeted Unweighted Perfect Rematching: Version II. Suppose P is a set of n people and H is a set of n houses, where we re-index the houses so that person p_i is currently matched with house h_i, and $B > 0$ is a positive integer. Associated to each person is a (possibly empty) subset of H, representing the set of houses that they prefer to their current house. Find the assignment f that maximizes *the number of people re-matched to a house they prefer to their current house minus the number of people re-matched to a house that they prefer less than their current house,* where $f : P \to H$ matches people to houses such that 1) f is a bijection, and 2) f assigns at most B people to a house they prefer less than their current house.

Note that the difference between Version I and Version II is simply whether the fraction of the compensation budget is charged against the market maker's profit, or whether it represents a fixed, sunk cost that cannot be recovered. In the example given above where the market maker has a small number B of new vacant luxury houses that everyone would prefer to their current house, Version I makes more sense: the preferred solution is the one that makes the maximum number of people happier: even if we assign all B of the new luxury homes, this is preferred to a solution that makes fewer people happier while leaving some of the new luxury homes vacant. On the other hand, if everyone who is re-matched to a hotel room they prefer to their current room pays the market maker $100, while everyone who we re-match to a hotel room that they like less than their current room needs to be compensated $100, then we clearly prefer the solution that maximizes the difference between the number of people who pay us and the number of people we must pay, so this is exactly modeled by Version II.

The above formulations presume that the value of each (positive or negative) re-matching is exactly the same, but in many problems, in addition to modeling a match that is more or less desired, we wish to quantify by how much the re-match is preferred or disliked. We can view the positive weight edge e_{ij} as the amount person i would pay to be re-matched to house j. Then the negative weight edges represent the amount it would cost to have person i give up his/her current house and instead accept house j, where we also have the additional constraint that the negative weight edges can sum to at most $-B$ (Again, we can consider a version I scenario where the market maker's profit is the total sum of the weights of the positive weight edges minus the sum of the negative weight edges of the re-matching, or a version II scenario, where the market maker has been required to put B dollars in escrow, and is now only trying to maximizes the sum of the weights of the positive weight edges).

Budgeted Perfect Rematching (Version II) Suppose P is a set of n people and H is a set of n houses, where we re-index the houses so that person p_i is currently matched with house h_i, and $B > 0$ is a positive integer. Let w_{ij} be the positive or negative weight associated with matching p_i to h_j with $w_{ii} = 0$, for all i. Find the assignment f that maximizes the sum of the weights of the edges in the assignment such that 1) f is a bijection, and 2) the negative weight edges sum to at most $-B$. (Note: for **Version I**, instead maximize the sum of the *positive* weight edges in the assignment).

Note that if there is unlimited budget, unweighted or weighted versions of reassignment are simply special cases of the ordinary assignment problem, and thus can be solved in polynomial time (we can transform a problem with positive, zero and negative weight edges to a problem with all positive weights by simply adding a fixed quantity equal to the magnitude of the maximum negative weight to all edges; the maximum solution to this assignment problem will be the same as to the original problem.) In this paper, we show that the unweighted budgeted reassignment problem can be solved in polynomial time, while the weighted budgeted reassignment problem is NP-hard. Using our new polynomial time algorithm for unweighted budgeted reassignment, we also present a set of simulations to begin to explore sets of preferences and distributions where allowing a small budget can unlock a lot of value for the market maker.

3 RELATED WORK

Our problem is most related to the "housing allocation with existing tenants" problem posed by [Abdulkadiroğlu and Sönmez, 1999]. However, rather than look at maximizing happiness, the authors are concerned with an analogy to stable matching and mechanisms that are Pareto optimal and strategy proof. Further related work in the strategy-proof setting can be found in [Miyagawa, 2001, Sönmez and Ünver, 2005, Svensson and Larsson, 2002]. A natural variant to this problem in the special case where dorm rooms are allocated when seniors graduate and freshmen have no existing allocation yet, but sophomores and juniors are already assigned, was studied by [Kurino, 2014].

More generally, the reassignment problem can be seen in the context of work that has been done in a variety of application domains to model *barter,* where barter is defined as the exchange of goods and services without the use of money. The most famous applications where barter has been studied are in domains in which money is repugnant, such as kidney exchange [Abraham et al., 2007, Roth et al., 2005]; this is also clearly in the same bipartite graph setting we are concerned with, i.e. we are also matching single, indivisible goods. However, our work does not apply to the kidney exchange domain because 1) we do not constrain the exchanges to occur along short cycles (as is required in the kidney exchange application, see for example [Abraham et al., 2007]) and 2) the focus of this paper are in applications where there is limited liquidity that can be injected into the market in a feasible and moral way (note that this does not have to be by use of money: for example, while school choice is often given as an example where monetary incentives are repugnant [Roth, 2007], but we can introduce limited liquidity without any form of monetary incentives if we allow the unbalancing/opening of additional seats in the most popular schools, see Section 7.)

In fact, from late medieval European trade fairs [Boerner and Hatfield, 2010, 2014] to distressed economies in the present, we find real-world instances where barter is preferred but a limited amount of money exists to augment exchange. In present day economies, such a scenario is in fact common amidst hyperinflation and liquidity crunches, as in the recent economic crisis in Zimbabwe in 2009 [Makochekanwa et al., 2009], as well as Greece in 2011 [Donadio, 2011]. Greece saw the emergence of online barter markets such as www.tradenow.com [Now, 2018]. While such markets might seem to represent edge cases, anthropological literature has illustrated that barter is a commonly observed social phenomenon in a number of informal markets. Examples include exchange practices amongst Andean farmers around Lake Titicaca [Orlove et al., 1986], as well as in Nepal [Humphrey, 1992]. A paper of Haddawy et al [Haddawy et al., 2005] in EC 2005 took a stab at modeling this problem, with multiple objects owned by each market participant, as a min-cost circulation problem. The scenario they consider is more general than that of this paper where each person has a single indivisible object to swap.

Many ecommerce portals currently exist that can be thought of as specialized barter markets. Websites that currently serve as barter portals include Read It Swap It for books [It, 2018], Intervac for holiday homes [Intervac, 2018], and the National Odd-Shoe Exchange for people with feet of two different sizes who wish to swap shoes instead of buying a pair for each foot [Exchange, 2018]. In any of these markets, a limited amount of cash or supplementary goods could be used to improve market clearing. We did not find a lot of previous theoretical work on modeling these types of online barter networks, apart from a nice paper of [Abbassi et al., 2013], which relate this problem to work on recommender systems, and a recent paper of Rappaz et al. [Rappaz et al., 2017].

We note that in the used goods market, agents who own particular goods might need to be incentivized by something like the reassignment market we envision in order to be willing to exchange goods even if this is in their rational best interest. This has been studied and is called "the endowment effect", the concept that individuals place a greater value on the things they own simply because they own them [Kahneman et al., 1990]. As Erickson and Fuster state in a recent review on the endowment effect,

> "By neglecting the endowment effect, people will tend to accumulate more belongings than optimal if storing things is costly." [Ericson and Fuster, 2013]

In other words, if the owner's value for the good is lower than the value she could get by selling it, then storing the good has a cost.

Genovese and Mayer showed that during a housing slump in the Boston condo market, owners who had bought their condos at higher prices spent more time holding onto their homes to get a good price that more closely matched their original purchase price [Genesove and Mayer, 2001]. An analogue of this is college students who accumulate expensive textbooks. Consider students that buy their textbooks new at the bookstore because they need the book immediately. They will make a greater loss when they sell their books used than the students that originally bought the book used. Because the textbook may have some future value, however minimal, students often choose to keep their textbooks for longer and postpone selling to see if a better buying price can be found later. Unfortunately, textbooks depreciate quickly with time, meaning that they often do not. Barter might be a means by which to overcome the endowment effect in the context of used goods.

While experiments are yet to be conducted on how agent behaviors differ given the medium of exchange, this hypothesis holds ground in empirical observation of the limitations of current used goods markets. Goods amenable to barter are clothes, which for many have seasonal value, and miscellaneous rare items that would otherwise not be on the market because of limited demand. Thus, it is worth exploring the role mixed monetary and non-monetary markets could play in the exchange of goods and services within communities. In the future, such explorations of mixing money and barter might allow us to create new types of markets and paradigms of exchange.

This paper is thus only a first foray into a potentially rich area with many open problems and applications; for expository purposes we mainly employ the language of "people" and "houses"; we discuss extensions to school choice in Section 7.

4 BUDGETED WEIGHTED PERFECT REMATCHING IS NP-HARD

Theorem 4.1. *Budgeted weighted perfect rematching is NP-hard.*

Proof Reduction from knapsack. Consider an instance of knapsack with objects $x_1, \ldots x_n$, and bound B, where object x_i has associated profit p_i and size s_i The goal is to select the set of objects whose profit is maximum, such that the sum of their sizes is $\leq B$. We transform this instance of knapsack to an instance of the budgeted weighted perfect rematching problem with budget B as follows: For each object x_i, we construct a set of two new people p_{yi} and p_{zi} and associated matched houses h_{yi} and h_{zi}, such that p_{yi} prefers h_{zi} to his current house with a weight of $p_i + s_i$ (for version I; for version II make the weight just p_i), and p_{zi} dislikes h_{yi} with a weight of $-s_i$ (both versions I and II). For every other house, that is not h_{yi} or h_{zi}, both p_{yi} and h_{zi} dislikes it with a weight $> B$ (so that the only re-matches that are allowed are between the pair). Then solutions to knapsack that include object x_i correspond to re-matchings that swap houses between p_{yi} and p_{zi}. So the reassignment problem with budget B has a solution with benefit $> X$ iff the corresponding instance of knapsack has a solution of profit $> X$. □

5 UNWEIGHTED BUDGETED PERFECT REMATCHING IN POLYNOMIAL TIME

Theorem 5.1. *Budgeted unweighted perfect rematching with budget B is solvable in polynomial time (for both version I and version II formulations.)*

Proof. We first consider the best perfect re-matching with the additional constraint that exactly b people are matched to houses they prefer less than their current house. Trying all integers b with $0 \leq b \leq B$ and returning the perfect re-matching of maximum benefit (where the maximum benefit is either the number of people re-matched to houses they prefer in each solution (Version I) or that number minus the up to b people who are made unhappy (Version II)) then solves our problem. Set positive $\epsilon < 1/n$. We form a bipartite graph $G_b = ((P \cup Q) \cup (H \cup G), E)$, where P and H correspond to people and houses as before (with the houses re-indexed so that p_i is currently matched to h_i), Q is a set of b new ghost people, and G is a set of b new ghost houses. We place an edge e_{ii} of weight ϵ between person p_i matched to house h_i; and for $i \neq j$ we place an edge $e_{ij} \in E$ of weight 1 whenever p_i prefers h_j to h_i. The new ghost people Q get weight ϵ edges to all the houses in H, and the new ghost houses G get weight ϵ edges to all the people P. All negative weight edges between G and H are removed. Note that there are *no* edges between the sets Q and G in this construction (we cannot match ghost people with ghost houses). Then any perfect matching in this graph matches b people with ghost houses (and b houses with ghost people). Return the perfect matching of maximum weight W_i in this bipartite graph G_b: it matches $n - b$ people with either their current house or a house they prefer to their current house. Now modify G_b as follows: Let P_u be the set of b original people matched to new ghost houses, and H_u be the set of original houses matched to ghost people. Match P_u back to H_u arbitrarily. First, this matching is within budget: everything but the b people are matched using edges of non-negative weight, so this matching costs at most b which is $\leq B$. Second, the solution to the max weight budgeted matching must appear among the G_bs. Suppose the max weight budgeting matching has exactly r negative edges. Then we return it when we consider G_r. □

Note that the cost of this algorithm is equivalent to solving B versions of the unweighted perfect rematching problem (and taking the maximum), so the running time is simply $O(Bn^{2.5} \log n)$. We note that several natural extensions to budgeted unweighted perfect matching can be solved using the same algorithm: we can handle variable positive weights on the edges in our call to weighted perfect matching, provided the negative weight edges are of constant weight and cost, i.e. the amount it costs us to compensate one of the up to B people who do not receive a house at least good as their own is the same for all people. In addition, we can handle a model with two classes of services: where up to B people from class II can be re-matched to a house they like less and compensated, but people from class I are guaranteed to never be assigned to a house they like less that their current house. This is easily modeled by only giving edges to the ghost houses from class II people. And it allows for situations where there are premium customers (such as airline flights).

6 SIMULATIONS

The existence of a polynomial time algorithm for unweighted budgeted rematching, makes it interesting to ask, for what distributions and types of sets of preferences that might arise in real-world examples, might allowing a small budget unlock substantial value for the market maker. We first define two simple extremes where it seems that having a budget greater than 1 does not typically help: in the *Random Preference Model* preferences among people are completely uncorrelated, whereas in the *Total Order Model* an objective order of desirability is agreed on by everyone. We next explore two additional scenarios (that we call the *Neighborhood Order Model* and the *Orthogonal Feature Model*) where we both argue that there exist market scenarios where the structure of preferences will plausibly resemble these models, and further, where allowing a budget can increase overall happiness in typical instances that arise in these models.

The Random Preference Model In this model, there is no correlation between the preferences of different individuals, i.e. each person independently chooses a random permutation of the objects as their own individual ranking of all objects. Once a global random perfect matching is chosen, each person is willing to be re-matched to any object that ranks higher in their own ranking than that of the object to which they were matched. This model is unlikely to make sense when matching houses, since many people will find the same houses objectively nicer or less nice, but might make sense when choosing artwork, for example.

The Total Order Model In this model, there is an objective sense in which objects are better and worse, so that each person has the same ranking of all the objects. Once a global random perfect matching is chosen, each person is willing to be re-matched to any object that ranks higher in their own ranking than that of the object to which they were matched. This model makes some sense when matching houses, but makes even more sense when assigning office space, where square footage and location might be universally agreed on measures of which offices are most desirable.

The Neighborhood Order Model In this model, houses are partitioned into sets called neighborhoods, and people are looking for houses only in their current neighborhood. Within their neighborhood, the houses follow the total order model, so that everyone agrees on what is the best house in the neighborhood.

The Orthogonal Feature Model In this model, houses are randomly assigned to one of three tiers: excellent, medium, and worst according to 3 separate characteristics or features: call them

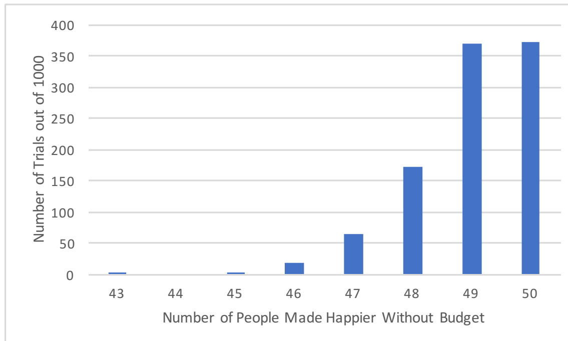

Fig. 1. Over a set of 1000 trials of the random preference model, histogram of the number of trials when n people are can be made happier (with a budget of 0).

(a,b,c). Then each person is randomly assigned one of the six permutations of abc: i.e. bac means they care most about characteristic b, then second most about characteristic a, then third most about characteristic c. People are initially randomly matched according to a global matching; they prefer any house that is better than theirs in the characteristic they prefer most, if there's a tie, then they consider the 2nd characteristic, if there's again a tie they go to the third characteristic. If some house is tied with their current house in all 3 characteristics, they prefer to stay in their current house rather than move. This models a more realistic scenario where houses can be objectively ordered on individual characteristics (size, number of bathrooms, curb appeal) but the individuals might set different levels of importance to having each of these features.

The above models only scratch the surface for what sorts of preference sets might be found in reassignment problems in real world applications. However, we find that they each behave quite differently in terms of how much more gain in happiness allowing a budget can unlock.

6.1 The random preference model

We claim that in the random preference model, it is very likely that we can make a great majority of the people happier, even without using any budget. People who certainly won't be made happier, however, are people who were already initially matched with their first-choice house: based on a variant of the famous Hat Check Problem [Scoville, 1966] as n gets large we expect no one to be matched with their top choice with probability $1/e$. We explored 1000 simulations of the random preference model using 50 people and 50 objects. Figure 1 gives a histogram of the number of people who can be made happier with a budget of 0. On the other hand, Figure 2 gives a histogram of the number of people who are matched to their top choice house. The similarity of distribution of these two histograms, leads us to conjecture the following (though we certainly do not have a proof):

Conjecture 6.1. *If houses and preferences are assigned according to the random preference model, once we remove the set of people who have been assigned to their top choice house and the houses they are matched with, the remaining graph has a perfect matching with high probability.*

Note that this conjecture is not obvious: in particular, people that are already assigned to their second choice house have only degree 1 in this graph. If we remove instead people assigned to their first and second choice houses, and their associated edges, then people have degree at least 2 in this graph. Alan Frieze proved in [Frieze, 2005] that the set of random bipartite graphs with $n + n$ vertices and cn random edges and minimum degree at least two, has a perfect matching w.h.p.: this

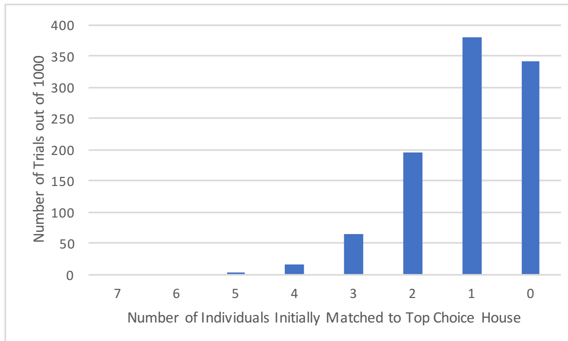

Fig. 2. The number of people initially matched with their top choice house over a set of 1000 trials of the random preference model.

is a different random model: plus, in our model, some house could have degree only 1 still (though this is not very likely, since half the people expect to have edges to at least half the houses).

Note, however, that if the above conjecture is true, that in most cases allowing a budget B does not improve the objective function for the random preference model.

6.2 The total order model and the neighborhood order model

In the total order model, there are no cycles, so with a budget of 0, there is no re-matching that makes anyone happier than just staying in their current house.

With a budget of 1, however, we can pay off the person in the top-ranked office to move to the bottom-ranked office, and then everyone else can move up by one rank, making everyone else happier. Clearly a budget of size > 1 conveys no further advantage, since everyone save that top-ranked person has already been made happier.

The neighborhood order model consists essentially of several copies of the total order model. In the neighborhood order model, we can assign a budget of 1 to each neighborhood, and do something similar to the total order model. So this is a fairly straightforward scenario where having a budget helps.

6.3 The orthogonal feature model

The orthogonal feature model begins to get at an interesting set of preferences that might successfully model many real-world problems. However, we no longer have any intuition on the best way to predict the results theoretically. Instead, we build our intuition through simulation: we simulated 50 different random matchings and preference graphs according to the orthogonal feature model using 50 people and 50 objects. Figure 3, for each one of our random instances of the orthogonal feature model, compares the maximum benefit obtained in Version II on the Y axis, as the budget B is increased along the X axis. As can be seen, there are instances where allowing up to 3 people to be made unhappy will increase overall happiness, but most of the time a budget of 1 or 2 suffices.

7 DISCUSSION

We have introduced variants of the reassignment problem, some of which are solvable using algorithms for weighted bipartite perfect matching, and some of which are computationally difficult. One natural open question is if the extra $O(B)$ factor in the running time of the budgeted unweighted algorithm is truly necessary. Another question is whether there are approximation algorithms for

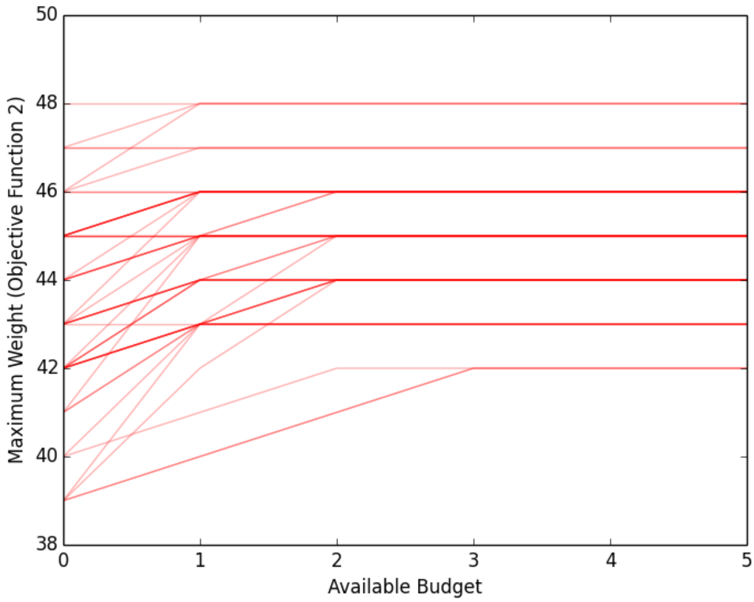

Fig. 3. Each line represents one of the 50 random instances of preferences we generated according to the orthogonal feature model. The Y axis shows the Version II benefit (the number of people we can make happier minus the number of people we make unhappier) that is achieved as the maximum available budget for that instance is increased along the X axis. It is highly variable when a budget helps and by how much.

versions of this problem that are NP-hard: where there has been a lot of work on approximate maximum matching lately (see [Duan and Pettie, 2014]), most of this work does not generalize to scenarios where one has to return an approximately optimal *perfect* matching. Perhaps the budget can help us have more flexibility to use approximate rather than exact matchings in some settings, since some limited number of unmatched vertices can be compensated instead of matched.

Another interesting open question is looking at the generalization of our one-to-one reassignment problem to the many-to-many problem of school choice. School choice has been extensively studied in the literature, but often in a stable or Pareto-optimal setting, rather than in a setting of attempting to maximize global happiness [Abdulkadiroğlu et al., 2005a,b, Abdulkadiroglu and Sönmez, 2003, Manlove, 2013] Here, we consider a city with a set of elementary schools that have a fixed number of seats, where there are enough seats between all the schools to seat all new Kindergarten students. Initial assignments to schools are by lottery, but based on neighborhood proximity or school characteristics, there might exist sets of families that would want to swap schools. How well could we do in satisfying family preference, particularly in the "budgeted" case where we would be allowed to satisfy additional families by opening a limited number of extra seats in the more desired schools?

In the case where a set of B additional seats is tolerated anywhere in the district, independent to how they are partitioned among the schools, then it is easy to see that we can handle this exactly as before: students have low weight edges to all the seats in the school to which they are currently assigned, and high weight edges to all the seats in the schools they prefer to their current school. However, it is much more realistic to assume a bound B on the number of extra seats that could be opened in *each* school, or even a bound b_1 of extra seats that could be opened in school 1, and a bound b_2 of extra seats that could be opened in school 2, where these bounds could allow different capacities for overflow at the various schools. If k, the number of schools is constant, given constraints $(b_1, \ldots, b_k)$ for the number of additional seats that each school tolerates, we can exhaustively create ghost seats, mimicking the construction of Theorem 5.1. The number of different k-tuples where B additional seats are places such that the number of additional seats in school i is at most b_i which is less than B, is upper-bounded by the number of k-tuples where the number of additional seats in school i is less than B, which is equal to the number of ways to place B identical balls into k distinct boxes, which is $((B + k - 1)$ choose $k - 1)$ which is $O(B^k)$ and polynomial in B when k is a constant. However, there might exist a much better algorithm that does not involve exhaustively checking all possible ways to allocate different possible numbers of extra seats to schools.

ACKNOWLEDGMENTS

We thank Christine Cheng for helpful discussions.

REFERENCES

Zeinab Abbassi, Laks VS Lakshmanan, and Min Xie. 2013. Fair recommendations for online barter exchange networks.. In *WebDB*. Citeseer, 43–48.

Atila Abdulkadiroğlu, Parag A Pathak, and Alvin E Roth. 2005a. The New York City high school match. *American Economic Review* (2005), 364–367.

Atila Abdulkadiroğlu, Parag A Pathak, Alvin E Roth, and Tayfun Sönmez. 2005b. The Boston public school match. *American Economic Review* (2005), 368–371.

Atila Abdulkadiroğlu and Tayfun Sönmez. 1999. House allocation with existing tenants. *Journal of Economic Theory* 88, 2 (1999), 233–260.

Atila Abdulkadiroglu and Tayfun Sönmez. 2003. School choice: A mechanism design approach. *The American Economic Review* 93, 3 (2003), 729–747.

David J Abraham, Avrim Blum, and Tuomas Sandholm. 2007. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce*. ACM, 295–304.

André Berger, Vincenzo Bonifaci, Fabrizio Grandoni, and Guido Schäfer. 2008. Budgeted matching and budgeted matroid intersection via the gasoline puzzle. In *International Conference on Integer Programming and Combinatorial Optimization*. Springer, 273–287.

Lars Boerner and John William Hatfield. 2010. The economics of debt clearing mechanisms. *Available at SSRN 1685149* (2010).

Lars Boerner and John William Hatfield. 2014. The design of debt clearing markets: clearinghouse mechanisms in pre-industrial Europe. *Available at SSRN 2183385* (2014).

Rachel Donadio. 2011. Battered by economic crisis, Greeks turn to barter networks. *The New York Times* 1 (2011).

Ran Duan and Seth Pettie. 2014. Linear-time approximation for maximum weight matching. *Journal of the ACM (JACM)* 61, 1 (2014), 1.

Keith M Marzilli Ericson and Andreas Fuster. 2013. *The endowment effect.* Technical Report. National Bureau of Economic Research.

National Odd Shoe Exchange. 2018. National Odd Shoe Exchange. (2018). http://www.oddshoe.org/

Alan Frieze. 2005. Perfect matchings in random bipartite graphs with minimal degree at least 2. *Random Structures & Algorithms* 26, 3 (2005), 319–358.

Harold N Gabow and Robert E Tarjan. 1991. Faster scaling algorithms for general graph matching problems. *Journal of the ACM (JACM)* 38, 4 (1991), 815–853.

David Genesove and Christopher Mayer. 2001. *Loss aversion and seller behavior: Evidence from the housing market.* Technical Report. National Bureau of Economic Research.

Peter Haddawy, Christine Cheng, Namthip Rujikeadkumjorn, and Khaimook Dhananaiyapergse. 2005. Optimizing ad hoc trade in a commercial barter trade exchange. *Electronic Commerce Research and Applications* 4, 4 (2005), 299–314.

Philip Hall. 1935. On representatives of subsets. *Journal of the London Mathematical Society* 1, 1 (1935), 26–30.

Caroline Humphrey. 1992. Fair dealing, just rewards: the ethics of barter in North-East Nepal. *Barter, exchange and value: an anthropological approach* (1992), 107–141.

Intervac. 2018. Intervac Home Exchange. (2018). https://us.intervac-homeexchange.com/

Read It Swap It. 2018. Read It Swap It. (2018). http://www.readitswapit.co.uk/TheLibrary.aspx

Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. 1990. Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy* (1990), 1325–1348.

Harold W Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics (NRL)* 2, 1-2 (1955), 83–97.

Morimitsu Kurino. 2014. House allocation with overlapping generations. *American Economic Journal: Microeconomics* 6, 1 (2014), 258–89.

László Lovász and Michael D Plummer. 2009. *Matching theory*. Vol. 367. American Mathematical Soc.

Albert Makochekanwa and others. 2009. Clothed in rags by hyperinflation: the case of Zimbabwe. *Munich Personal RePEc Archive. Paper* 28863 (2009).

David F Manlove. 2013. *Algorithmics of matching under preferences*. Vol. 2. World Scientific.

Eiichi Miyagawa. 2001. House allocation with transfers. *Journal of Economic Theory* 100, 2 (2001), 329–355.

James Munkres. 1957. Algorithms for the assignment and transportation problems. *J. Soc. Indust. Appl. Math.* 5, 1 (1957), 32–38.

Trade Now. 2018. Trade Now. (2018). http://tradenow.gr/en/welcome.aspx

James B Orlin and Ravindra K Ahuja. 1992. New scaling algorithms for the assignment and minimum mean cycle problems. *Mathematical Programming* 54, 1-3 (1992), 41–56.

Benjamin S Orlove, James M Acheson, John Clammer, Thomas Crump, Stephen Gudeman, David Guillet, Olivia Harris, and Henry J Rutz. 1986. Barter and cash sale on Lake Titicaca: A test of competing approaches [and comments and replies]. *Current Anthropology* (1986), 85–106.

Jérémie Rappaz, Maria-Luiza Vladarean, Julian McAuley, and Michele Catasta. 2017. Bartering books to beers: a recommender system for exchange platforms. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM, 505–514.

Alvin E Roth. 2007. Repugnance as a constraint on markets. *Journal of Economic Perspectives* 21, 3 (2007), 37–58.

Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. 2005. Pairwise kidney exchange. *Journal of Economic Theory* 125, 2 (2005), 151–188.

Richard Scoville. 1966. The hat-check problem. *The American Mathematical Monthly* 73, 3 (1966), 262–265.

Tayfun Sönmez and M Utku Ünver. 2005. House allocation with existing tenants: an equivalence. *Games and Economic Behavior* 52, 1 (2005), 153–185.

Lars-Gunnar Svensson and Bo Larsson. 2002. Strategy-proof and nonbossy allocation of indivisible goods and money. *Economic Theory* 20, 3 (2002), 483–502.