



Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejorComplexity of finding Pareto-efficient allocations of highest welfare[☆]Péter Biró^{a,b}, Jens Gudmundsson^{c,*}^aInstitute of Economics, Research Centre for Economic and Regional Studies, Hungarian Academy of Sciences, Budapest, Hungary^bDepartment of Operations Research and Actuarial Sciences, Corvinus University of Budapest, Hungary^cDepartment of Food and Resource Economics, University of Copenhagen, Denmark

ARTICLE INFO

Article history:

Received 30 January 2019

Accepted 4 March 2020

Available online xxx

Keywords:

Assignment

Pareto-efficiency

Welfare-maximization

Complexity

Integer programming

ABSTRACT

We allocate objects to agents as exemplified primarily by school choice. Welfare judgments of the object-allocating agency are encoded as edge weights in the acceptability graph. The welfare of an allocation is the sum of its edge weights. We introduce the *constrained welfare-maximizing solution*, which is the allocation of highest welfare among the Pareto-efficient allocations. We identify conditions under which this solution is easily determined from a computational point of view. For the unrestricted case, we formulate an integer program and find this to be viable in practice as it quickly solves a real-world instance of kindergarten allocation and large-scale simulated instances. Incentives to report preferences truthfully are discussed briefly.

© 2020 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

We consider the allocation of objects to agents, such as school seats to students, in the absence of monetary transfers. We take as given that the assignment should reflect the agents' preferences and operationalize this by restricting to Pareto-efficient allocations.¹ Typically, not all agents can receive their first choice and the more popular objects have to be rationed. How they are rationed reflects a welfare judgment on behalf of the object-allocating agency, say in terms of what is fair or socially optimal. For instance, costs of transportation may prohibit admitting a stu-

dent to a highly preferred but remote school or there may be benefits to sending children on the same street to the same kindergarten even though the families' preferences differ. We model this in a simple yet surprisingly flexible way: assigning agent i object a creates welfare $w(i, a)$, and the welfare of an allocation is the sum of these terms. We use tools of economics, computer science, and operational research to address the following questions: Under which conditions does a Pareto-efficient allocation maximize welfare? When there is a trade-off between Pareto-efficiency and welfare-maximization, how should the objects be allocated? And finally, given that the problem can encompass a large number of agents in practice, can we find a desirable solution efficiently from a computational point of view?

In addressing these questions, we refer for the most part to school choice (Abdulkadiroğlu and Sönmez, 2003; Shi, 2016), but the problem extends to a wide range of applications, some well-known and some new within the field of OR. These include university admission, resident allocation (Bronfman, Hassidim, Afek, Romm, Shreberk, Hassidim, & Massler, 2015), dormitory room allocation (Perach, Polak, & Rothblum, 2008), deceased organ donation, social housing, and refugee allocation (Andersson & Ehlers, 2017; Moraga & Rapoport, 2014; Delacrétaz, Kominers, & Teytelboym, 2016; Trapp, Teytelboym, Martinello, Andersson, & Ahani, 2018). Present in each of these applications is a centralized object-allocating institution (a "planner") with its own objective function that should be taken into consideration in parallel with the agents' preferences. We refer to Section 2 for more detailed examples on

^{*} We are grateful to the anonymous reviewers for their constructive suggestions. We thank participants at the Lisbon Meetings (2017), the Conference on Mechanism and Institution Design (Durham, 2018), Frontiers of Market Design (Lund, 2018), Matching in Practice (Mannheim, 2018), the EAADS Workshop (Kosice, 2018), the Conference on School Choice and Reform (Lisbon, 2019), the Conference on Economics Design (Budapest, 2019), and the seminar participants at Corvinus University for valuable comments. Péter Biró acknowledges the support of the Hungarian Academy of Sciences under its Momentum Programme (LP2016-3/2018) and Cooperation of Excellences Grant (KEP-6/2018), and the Hungarian Scientific Research Fund, OTKA, Grant No. K128611. Jens Gudmundsson acknowledges the financial support of the Jan Wallander and Tom Hedelius Foundation.

^{*} Corresponding author.

E-mail addresses: peter.biro@krtk.mta.hu (P. Biró), jg@ifro.ku.dk (J. Gudmundsson).

¹ An allocation is Pareto-efficient if no other allocation leaves each agent at least as well off and some agent better off. As an example, if we order the agents and let them sequentially select their preferred object (among those that remain), then the final allocation will be Pareto-efficient. This procedure is known as *Serial Dictatorship*.

<https://doi.org/10.1016/j.ejor.2020.03.018>

0377-2217/© 2020 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

how the welfare levels w (the “edge weights”) can be set in order to cover a wide variety of objectives.

For the particular application of school choice, we propose an alternative to the priority-based approach that is most used in practice (see Section 1.1). Essentially, whereas our edge weights can be encoded with cardinal information, the ordinal priorities cannot.² If one student is prioritized over another, then this may show that the former lives closer to the school—but not how close, or how much closer. To illustrate further, say two students prefer school a to b and that student 1 lives next to b but also closest to (but some distance away from) a . An allocation is then Pareto-efficient as long as both students are assigned a school. By respecting the distance-based priorities, student 1 is admitted to a , student 2 to b , and both students require transportation. In contrast, by setting the edge weight between 1 and b high, our optimal solution will swap the assignment and thus reduces transportation costs while retaining Pareto-efficiency.

For our first result, we rely on a recent finding by Saban and Sethuraman (2015). They study the complexity of determining the outcome obtained by *Random Serial Dictatorship* and derive, as a byproduct, results on the decision problem termed SD FEASIBILITY. This asks: For a given profile of preferences, an agent i , and an object a , does there exist a serial dictatorship that assigns a to i ? Saban and Sethuraman (2015) show that SD FEASIBILITY is NP-complete even in a restricted environment.³ Using this result, our Theorem 1 shows that deciding whether there exists a Pareto-efficient, welfare-maximizing allocation is NP-complete with the same restrictions imposed as in the result of Saban and Sethuraman (2015). For these hard problems, relaxing the restrictions only makes the problems yet harder: they remain NP-hard, but the results get weaker.

Unless the preferences or the edge weights take on a particular form, there is little reason to believe that there actually exists a Pareto-efficient allocation that maximizes welfare. When there is a conflict between these desiderata, we propose to select a *constrained welfare-maximizing* allocation. This is a Pareto-efficient allocation of highest welfare among the Pareto-efficient allocations. We label the problem of finding such an allocation CONstrainedWELFAREMAX. This is harder than deciding whether there exists a Pareto-efficient, welfare-maximizing allocation, so CONstrainedWELFAREMAX is computationally tractable only under yet stronger conditions. As a first step, we restrict attention to *object-based weights*. Such weights $w(i, a)$ depend only on the object a . For school choice, this can be interpreted as the planner promoting a particular school or topic of study. In Theorem 2, we show that CONstrainedWELFAREMAX is NP-hard even under object-based weights and complete preferences.

The next result pertains to the case in which all agents rank the objects acceptable to them in the same way. This is a plausible restriction when there is an objective measure of quality on the objects, such as some schools providing objectively better education than others. If all agents rank the objects acceptable to them in the same way, then the condition of *common preferences* is satisfied. Theorem 3 shows that deciding whether there exists a Pareto-efficient, welfare-maximizing allocation is NP-complete even for balanced problems with object-based weights and common preferences. In addition, we derive a result that complements (Saban & Sethuraman, 2015) on SD FEASIBILITY.

² The same argument applies to preference intensity, which again is cardinal information that cannot be encoded in ordinal preferences. On this topic, Abdulkadiroğlu, Che, and Yasuda (2015) introduce *Choice-Augmented Deferred Acceptance*, which allows agents to express richer preference information. In particular, their agents report both a ranking over schools and a “target” school.

³ Specifically, it is NP-complete when agents have complete and strict preferences and the problem is “balanced” with an equal number of agents and objects.

We then proceed to identify computationally tractable cases. Theorem 4 shows that, for balanced problems with object-based weights and complete preferences, all serial dictatorships yield constrained welfare-maximizing allocations. Theorem 5 shows that CONstrainedWELFAREMAX is polynomial-time solvable under common and complete preferences using (Kuhn, 1955) Hungarian method. There are other ways of combining the four conditions introduced thus far, but the remaining cases can all be inferred from Theorem 1 through 5 as summarized in Fig. 1.

We introduce three additional conditions, each on its own is strong enough to make CONstrainedWELFAREMAX tractable. The first of these restricts to dichotomous preferences in which all agents are indifferent between all objects acceptable to them. This domain restriction is relevant when all objects are similar in quality. Under this condition, an allocation is Pareto-efficient if and only if it is of maximum cardinality. Theorem 6 shows that, under dichotomous preferences, a constrained welfare-maximizing allocation can be found using the Hungarian method.

The next condition is *aligned interests* and implies that assigning an agent a more preferred object leads to higher welfare. That is, the interests of the planner, to assign higher-valued objects, is aligned with the interests of the agents, to be assigned more preferred objects. Theorem 7 shows that, under aligned interests, each welfare-maximizing allocation is Pareto-efficient. In consequence, there exists a Pareto-efficient, welfare-maximizing allocation, and we find it efficiently using the Hungarian method.

The final restriction is to *agent-based weights*, which are such that the weight $w(i, a)$ only depends on the agent i . Such weights are plausible for instance in merit-based university admissions: the objective is to admit the students with the highest grades, but it is less important where they are admitted. Theorem 8 shows that, under agent-based weights, there exists a Pareto-efficient, welfare-maximizing allocation. Again, we find it efficiently using the Hungarian method. We emphasize also that, for each of the positive results of Theorem 4 through 8, we permit indifference in the preferences, whereas the hardness results of Theorems 1, 2, and 3 are obtained when restricting to strict preferences. Generalizing by permitting indifference only makes these problems yet harder. Therefore, our polynomial-time algorithms remain efficient for strict preferences while our NP-hardness results still hold when we allow indifferences.

Moving away from the computational aspects, Theorem 9 provides a still related finding: under positive and object-based weights, each constrained welfare-maximizing allocation is of maximum cardinality. That is, there exists no way of allocating more objects even if the constraints of Pareto-efficiency and welfare-maximization are removed.

We also formulate an integer program (IP) to solve the general, unrestricted problem. To do so, we use a novel characterization, in the form of linear constraints, of Pareto-efficiency in the presence of preference ties. Theorem 10 provides a link between Pareto-efficiency and competitive equilibrium through a variation of the well-known second welfare theorem. Theorem 11 summarizes the IP formulation.

Finally, we consider the strategic properties of our solution concept. Theorem 12 establishes a positive result when the problem is sufficiently restricted. That is, with some conditions on the preferences and the weights, we can select constrained welfare-maximizing allocations in a way that incentivizes the agents to report preferences truthfully. However, we also show that relaxing either of the conditions may allow for manipulation.

The paper is structured as follows. Next, we describe the related literature. We introduce the model in Section 2 together with a series of examples and applications. In Section 3, we examine the complexity of finding constrained welfare-maximizing allocations, first through hardness results and then through tractable cases.

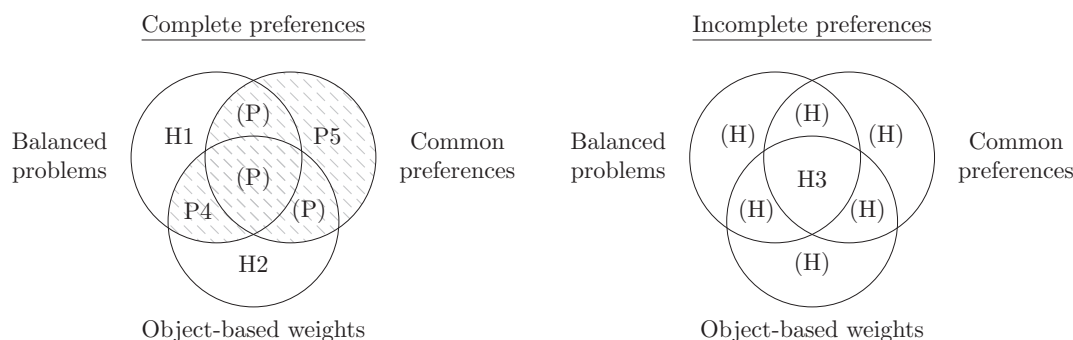


Fig. 1. Complexity of CONSTRAINEDWELFAREMAX. Numbers refer to theorems, “H” to NP-hardness, and “P” to polynomial-time solvable cases. Parentheses indicate that the case is covered by a stronger result.

In Section 4, we formulate the IP to solve CONSTRAINEDWELFAREMAX. In Section 5, we examine the 2016 kindergarten allocation in Harku, Estonia, with a focus on comparing the constrained welfare-maximizing allocation with other solutions. We discuss incentives issues in Section 6. We conclude in Section 7. Appendix A defines all solutions referred to throughout the paper more formally. Finally, Appendix B contains a simulation study that serves both to contrast the different solutions and to show that the integer programming approach is viable in solving larger problems quickly.

1.1. Related literature

A focal point in the literature are the *stable* (also known as justified envy-free and non-wasteful) allocations. Stable allocations refuse an agent a preferred object only if the object is assigned a higher-priority agent. Such an allocation can be computed through the *Deferred Acceptance* algorithm (DA; Gale & Shapley, 1962), which may be adapted to handle priority ties (Erdil & Ergin, 2008).⁴ These solutions have a particular structure as shown through the so called Rural Hospitals’ theorem (Roth, 1984a; 1986; Gale & Sotomayor, 1985). First, the same students are allocated in every stable allocation. Second, a school that fails to fill its seats at one stable allocation is assigned the same students at every stable allocation. Nowadays, DA may be the most used procedure in college admission and school choice programs around the world (for specific cases, see Abdulkadiroğlu, Pathak, & Roth, 2005a; Abdulkadiroğlu, Pathak, Roth, & Sönmez, 2005b; for a recent survey, see Biró, 2017). Not only is DA stable, but it also selects the student-optimal allocation among the stable allocations and it gives students incentives to report preferences truthfully. A recent development is to allow agents to report preferences only over a well-designed “menu” of schools (Ashlagi & Shi, 2016; Shi, 2015). In this way, the outcome of DA becomes closer to optimal from the point of view of the planner (in this case, the city of Boston) in as far as cutting down busing costs. Taken to its extreme, if the menus only contain a single school, then the planner can implement any allocation.

The second leading procedure is *Top Trading Cycles* (TTC; Shapley & Scarf, 1974; Abdulkadiroğlu and Sönmez, 2003), for instance used in New Orleans (Abdulkadiroğlu, Che, Pathak, Roth, & Tercieux, 2017). Like DA, TTC cannot be manipulated, but in contrast to DA, TTC is Pareto-efficient but not stable. *Serial Dictatorship* is another non-manipulable way of achieving a Pareto-efficient

allocation, used in Amsterdam’s school choice (de Haan, Gautier, Oosterbeek, & Van der Klaauw, 2018) and for residence allocation in Israel until 2014 (Bronfman et al., 2015). It is equivalent to TTC when schools share priorities. Another Pareto-efficient solution is *Immediate Acceptance*, which was used in Boston (Abdulkadiroğlu and Sönmez, 2003) and is still used in many applications. Its manipulability is considered its main issue, although it has still some desirable properties in regards to the expected utilitarian welfare (Abdulkadiroğlu, Che, & Yasuda, 2011). Finally, simple first-come first-served systems are sometimes used for course allocation, for example in almost every university in Hungary. Fig. 2 summarizes these approaches; see also Appendix A for more formal definitions.

The model of matching with contracts is an important extension that is well-studied for two-sided matching problems such as the match of doctors to hospital residency programs (see, for instance, Cechlárová & Fleiner, 2005; Fleiner, 2003; Hatfield & Milgrom, 2005). The extension is meaningful also in our allocation setting as it allows to assign objects to agents under different contractual terms. To illustrate, consider the Hungarian college admissions. Most programs can be attended under two possible contracts: either the student pays a fee or the state finances the studies (see Biró, 2011). There is still a trade-off as stricter rules apply to the state-funded contracts: the student has to graduate in a certain time and work in Hungary for some time thereafter. If she fails to meet these requirements, she has to pay back the funding with interest. Therefore, some students prefer to pay the tuition fee over taking part in the “free” state-funded programs (see also Shorrer & Sívágó, 2018). For our purposes, contracts are easy to include in the model but actually add very little. This is due to the strong implication of Pareto-efficiency: each agent is assigned her object under her most preferred contract. Otherwise, changing the terms of her contract is a Pareto-improvement. In particular, if we, for each agent and object, remove all contracts but the preferred one, then the set of Pareto-efficient allocations is unchanged. Thus, finding a constrained welfare-maximizing allocation in a model with contracts is no more difficult than finding one when there is just a single contract for each agent-object pair, which is equivalent to the simpler model without contracts.

A new application of operational research, receiving considerable attention following the 2015 European migrant crisis, is refugee allocation. While the larger problem of assigning refugees to countries is likely to be resolved using other criteria (Bansak, Hainmueller & Hangartner, 2017; Moraga and Rapoport, 2014), the assignment *within* countries can be viewed as an allocation problem in line with our stylized model. Specifically, we may think of the refugees as our agents and of the various locations that they can be resettled to as our objects (Delacrétaz et al., 2016) or turn things around and view citizens looking to host refugees as agents while treating the refugees as objects (Andersson & Ehlers, 2017).

⁴ The problem with ties is quite different from the strict case and there are several different notions of stability, such as weak, strong, and super stability. In the economics literature, focus has mainly been on weak stability. Weakly stable allocations always exist, but they may differ in size. Moreover, the problem of finding a weakly stable allocation of maximum size is NP-hard (Manlove, Irving, Iwama, Miyazaki, & Morita, 2002).




Fig. 2. The most common ways of taking preferences into account is by selecting a stable or a Pareto-efficient allocation. In parallel, welfare factors can be accounted for through priorities, or, as we are proposing, a linear welfare function. Here, “CWM” denotes the constrained welfare-maximizing solution.

Our edge weights can then encode, as Bansak, Ferwerda, Hainmueller, Dillon, Hangartner, Lawrence, and Weinstein (2018) put it, that “there are synergies between places and people” and that “certain characteristics will make a refugee a better match for a particular location”. For instance, the edge weights can be determined using the data-driven algorithm provided by Bansak et al. (2018) to represent the probability that a refugee will find employment within a location (see also Mossaad, Ferwerda, Lawrence, Weinstein, & Hainmueller, 2018; Bansak, Hainmueller, & Hangartner, 2016; Trapp et al., 2018). Once the infrastructure is in place to allow refugees to express their preferences in a safe and credible way, our IP can readily be used to find a welfare-maximizing refugee resettlement.

Lastly, we summarize some of the existing complexity results. Abraham, Cechlárová, Manlove, and Mehlhorn (2005) provide the first NP-hardness results by showing that finding a minimum size Pareto-efficient allocation is computationally hard. For two-sided matching, Irving, Leather, and Gusfield (1987) showed that finding an optimal stable matching for a linear welfare function is a tractable problem. In contrast, if the objective is to find an optimal allocation (say with respect to size or welfare) that is “as stable as possible” in that it minimizes the number of blocking pairs, then the problem is NP-hard (Biró, Manlove, & Mittal, 2010). We refer to Manlove (2013) for a comprehensive survey of related complexity results.

Regarding the optimization techniques used to tackle the above described computationally hard cases, integer programs have received significant attention in recent years. They have been used in more general settings such as allocation of papers to reviewers (Garg, Kavitha, Kumar, Mehlhorn, & Mestre, 2010) or for course allocation (Othman, Sandholm, & Budish, 2010), later implemented at Wharton College (Budish, Cachon, Kessler, & Othman, 2016). Further examples include the resident allocation problem with couples motivated by the US and Scottish applications (Biró, Manlove, & McBride, 2014) and the college admission problem with lower and common quotas (Ágoston, Biró, & McBride, 2016). Other OR techniques, based on Scarf’s lemma (Scarf, 1967), have also been proposed for the problem of matching with couples (Biró, Fleiner, & Irving, 2016) and in other many-to-one stable matching settings such as Nguyen, Nguyen, and Teytelboym (2019).

There are different reasons for why integer programming techniques had not been used to solve these problems until recently. First, the heuristic algorithms based on DA perform relatively well in practice, as illustrated by the US resident allocation program with couples (Roth & Peranson, 1999). Secondly, the problem sizes are relatively large (around 40,000 residents in the US, and 100,000 students in the Hungarian higher education admission scheme), and this large input size can be challenging for the IP solvers. However, new studies within computer science and OR show that even such large problems can be tractable with the IP approach (see, for instance, Firat, Briskorn, & Laugier, 2016). As an example, the NP-hard problem of having lower quotas for university programs was solved for a real 2008 instance of the Hungarian college admissions after a careful preprocessing and by using advanced IP techniques (Ágoston et al., 2016). We believe that our IP formulation for a special problem setting can be equally useful in




Fig. 3. Students and schools for Example 1. Edge weights represent distances.

practice and serve as a starting point when considering more general problems.

2. Model

There is a finite set of agents $N = \{1, 2, \dots\}$ and objects $A = \{a, b, \dots\}$. A problem is *balanced* if there are as many agents as there are objects. Agent i has *preference* R_i over objects acceptable to her, $A_i \subseteq A$, and not being assigned an object, \emptyset . She finds a at least as good as b whenever $a R_i b$. The strict relation is denoted P_i and the indifference relation I_i . For each $a \in A_i$, $a P_i \emptyset$. The bipartite *acceptability graph* $(N \cup A, E)$ has an edge $(i, a) \in E$ whenever a is acceptable to i . An *allocation* $x \subseteq E$ is an independent edge set (matching). If $(i, a) \in x$, then i is assigned object $x_i = a$. If i is not assigned an object, then $x_i = \emptyset$. The set of allocations is X . An allocation is *Pareto-efficient* if no allocation leaves each agent at least as well off and some agent better off. Thus, $x \in X$ is Pareto-efficient if there is no $y \in X$ such that, for each $i \in N$, $y_i R_i x_i$, and, for some $j \in N$, $y_j P_j x_j$. The social welfare of assigning agent i object $a \in A_i$ is represented by the *weight* $w(i, a) \geq 0$ on the edge $(i, a) \in E$. The *welfare* of $x \in X$ is $W(x) = \sum_{(i,a) \in x} w(i, a)$. An allocation $x \in X$ is *welfare-maximizing* if it creates the highest welfare among all allocations: for each $y \in X$, $W(x) \geq W(y)$. A Pareto-efficient allocation is *constrained welfare-maximizing* if it creates the highest welfare among the Pareto-efficient allocations.

Next, we illustrate the model through two examples in the context of school choice.

Example 1 (Distance-based school choice). In a school choice problem, the agents are students and the objects are seats at schools. A school with several seats is thus treated as several distinct objects which we assume students to be indifferent between. Typically, the student assignment has to respect some priorities, say to favor students closer to the school. We offer a different implementation of this through the edge weights inspired by a recent court case in Lund, Sweden (see Andersson, 2017, “Perceived issues”).

More students top-ranked school a than it had seats for, so its seats were assigned to the students living closest. Student 1 had 2.2 kilometers walking distance to her preferred school a and 3.2 kilometer to her assigned school b . Hence, the walking distance for 1 increased by one kilometer when she was placed at b rather than at a . In contrast, student 2 was assigned to a but would only have had to walk 250 meters further had she been placed at b . See Fig. 3 for an illustration.

The parents argued that student 1 should be given higher priority at a as 1 would lose more from being placed elsewhere. Two courts (Förvaltningsrätten and Kammarrätten) have ruled in favor




Fig. 4. Students and schools in Example 2. An edge weight of 1 represents walk-zone priority. For the sake of readability, we have left out the remaining zero-weight edges.

of the parents. One interpretation is that the judges considered the rejection of 1 at a fair from the perspective of the distance-based priorities, but that the allocation suggested by the parents was better from the planner's perspective as it reduced the total travel distance.

Example 2 (Walk zones in school choice). School choice priorities are typically coarse with many ties. For instance, all students within the school's walk zone may be equally prioritized. As an example, say that there are four students and four schools, each with a single seat. Each student prefers school a to b to c to d , so any complete allocation is Pareto-efficient. Walk-zone priority is given to student 1 at schools a and b , to 2 at c , and to 3 at d ; see Fig. 4.

Before running an algorithm like *Deferred Acceptance* or *Top Trading Cycles*, the priority ties are often broken randomly. This may give student 2 priority at school b (over 3 and 4) and 3 priority at c . Then both algorithms select the allocation $\{(1, a), (2, b), (3, c), (4, d)\}$, which sends only student 1 to a walk-zone school. On the other hand, if we use the edge weights to indicate whether the student is within the school's walk zone, then the allocation $\{(1, b), (2, c), (3, d), (4, a)\}$ is welfare-maximizing. It assigns all but student 4 to a walk-zone school. Hence, this reduces transportation costs while retaining Pareto-efficiency.

To finish this section, we want to stress that we can take many different objectives into account through the edge weights. In Example 1, we showed how one can minimize total walking distance, while we in the introduction discussed having weights represent the probabilities that refugees find employment at different locations. Next, we will see that many more objectives can be achieved through our solution concept.

Maximizing the number of agents assigned acceptable objects. As in Example 2, we may want to maximize the number of agents assigned to suitable places. Further examples are refugee allocation (Andersson & Ehlers, 2017), kindergarten allocation (Veski, Biró, Pöder, & Lauri, 2017), and timetable scheduling. In our model, this objective can be achieved through uniform weights (see Theorem 9).

Assigning important agents or positions. We may want to guarantee the allocation of a particularly important group of agents or objects. A practical example is (re-)allocation in the US Navy (Yang, Giampapa, & Sycara, 2003). This objective can be achieved by putting large weights on the edges incident with the important nodes in the graph (compare the set A° in Theorem 5).

Reallocation with initially assigned objects. In many applications agents either own or are initially assigned objects and the task is to reallocate the objects in a desirable way. Two examples are teacher reallocation in France (Combe, Tercieux, & Terrier, 2018) and kidney exchange (Roth, Sönmez, & Ünver, 2004). In these settings, one typically has to ensure that no agent receives an object less preferred than the one she is initially assigned. This is achieved in our model by putting large weights on the edges linking the agents to their initial assignments and their more preferred objects.

Affirmative action in school choice by minority reserves. In many school choice and college admission systems, there are distributional goals over the composition of the students. This is some-

times enforced by so-called minority reserves, meaning that students with a particular background or ethnicity have priority for some school seats. Real-world examples include school choice in the US (Abdulkadiroğlu and Sönmez, 2003) and college admissions in India (Aygün & Turhan, 2017; Sönmez & Yenmez, 2019) and Brazil (Aygün & Bó, 2017). This objective can be achieved in our model by putting large weights on the edges linking minority students to the school seats reserved for the minority students. As an example, a school with 100 seats that wants to admit 20 students with a particular socio-economic background can link 20 copies of its seats with high-weight edges to students qualifying for the affirmative action policy.

3. Complexity of constrained welfare-maximization

In this section, we study the complexity of finding a constrained welfare-maximizing allocation. We denote this problem **CONSTRAINEDWELFAREMAX**, and we show that it is NP-hard even under strong conditions. However, we also identify restrictions under which the problem is tractable. Throughout, all hardness results restrict to strict preferences whereas the positive results permit indifference. Moreover, the hard problems are stated in their most restricted form and imply NP-hardness for all less restricted settings.

3.1. Hardness results

For our first result, suppose that a single edge has non-zero weight. That is, there is an agent i and an object a such that $w(i, a) > 0$ and otherwise the weights are zero. An allocation is then welfare-maximizing if and only if it assigns a to i . Furthermore, say preferences are strict and complete ($A_i = A$); an allocation is then Pareto-efficient if and only if it is obtained through a serial dictatorship (Abdulkadiroğlu and Sönmez, 1998). Hence, there exists a Pareto-efficient, welfare-maximizing allocation if and only if there exists a serial dictatorship that assigns a to i . This problem, deciding whether such a serial dictatorship exists, is known as **SD FEASIBILITY** and shown by Saban and Sethuraman (2015), Theorem 2) to be NP-complete for balanced problems. Theorem 1 follows immediately.

Theorem 1. *Deciding whether there exists a Pareto-efficient, welfare-maximizing allocation is NP-complete even for balanced problems with complete preferences in which a single edge has non-zero weight.*

Proof. Checking Pareto-efficiency (Manlove, 2013, Section 6.2.2.1) and welfare-maximality of an allocation (the Hungarian method of Kuhn, 1955) can be done in polynomial time (even for unbalanced problems with indifferences and incomplete preferences). Hence, the problem is in NP. To show NP-hardness of our problem, we reduce from **SD FEASIBILITY** as described above. \square

Corollary 1 is an immediate implication of Theorem 1.

Corollary 1. *CONSTRAINEDWELFAREMAX is NP-hard even for balanced problems with complete preferences.*

Proof. By contradiction, if we easily could find a constrained welfare-maximizing allocation, then we just need to check whether it assigns a to i . If it does, then there exists a Pareto-efficient, welfare-maximizing allocation. If it does not, then no such allocation exists. But Theorem 1 shows that this is NP-complete. \square

This construction shows that a polynomial-time approximation algorithm that finds a "good" but not necessarily constrained welfare-maximizing allocation cannot give a meaningful worst-case welfare-guarantee. That is, the welfare-difference between the constrained welfare-maximizing allocation and a second-best allocation can be arbitrarily large, $w(i, a) = 0$. Indeed, neither of the

cases shown to be NP-hard in Theorems 1, 2 and 3 permits a meaningful polynomial-time approximation.⁵

We proceed to derive related results by imposing different restrictions. Generally, edge weights can depend on both the agent and the object, but for the upcoming results it suffices to restrict attention to object-based weights. That is, the social welfare of assigning an object to an agent depends only on the object. Theorem 2 shows that the NP-hardness result extends to the case of object-based weights and complete preferences.

Definition 1 (Object-based weights). For each $\{i, j\} \subseteq N$ and $a \in A_i \cap A_j$, $w(i, a) = w(j, a)$.

Theorem 2. *CONSTRAINEDWELFAREMAX is NP-hard even under object-based weights and complete preferences.*

Proof. We reduce from the problem of assigning objects in a Pareto-efficient way but to as few agents as possible. Finding such a minimum size Pareto-efficient allocation is NP-hard (Abraham et al., 2005).

Extend an arbitrary instance $\mathcal{I} = (N, A, E, R, w)$ to $\mathcal{I}^* = (N, A \cup A^*, E^*, R^*, w^*)$ such that the properties required in the statement of Theorem 2 are satisfied for \mathcal{I}^* . In particular,

- Add objects A^* such that $|A^*| = |N|$;
- Complete the preferences through a complete acceptability graph: $E^* = N \times (A \cup A^*)$;
- New preference R_i^* extends R_i : the acceptable objects A_i at \mathcal{I} are in the same order at the top of R_i^* , followed by A^* in any order, followed by the unacceptable objects of \mathcal{I} ;
- The object-based weights are $w^*(i, a) = 0$ for $a \in A$ and $w^*(i, a) = 1$ for $a \in A^*$.

In this way, \mathcal{I}^* is guaranteed to have at least as many objects as agents.

Let x^* be a constrained welfare-maximizing allocation in \mathcal{I}^* . By Pareto-efficiency, as each of the $|N|$ agents i prefers at least $|A^*| = |N|$ objects in \mathcal{I}^* to those unacceptable in \mathcal{I} , $x_i^* \in A_i \cup A^*$. Define the corresponding allocation x in \mathcal{I} such that $x_i = x_i^*$ if $x_i^* \in A_i$ and $x_i = \emptyset$ if $x_i^* \in A^*$. As x^* is Pareto-efficient in \mathcal{I}^* , x is Pareto-efficient in \mathcal{I} . The welfare $W(x^*)$ is the number of assigned A^* -object in x^* , so the number of unassigned agents in x . Therefore, if we can efficiently find the constrained welfare-maximizing x^* in \mathcal{I}^* —that is, a Pareto-efficient allocation that assigns the most A^* -objects—then we can efficiently find a minimum size Pareto-efficient allocation in \mathcal{I} . But this is NP-hard. \square

In some applications, there may exist an objective ranking of the objects (Alpern & Katrantzi, 2009). That is, agents' preferences are derived from a common preference \succsim on A and whenever an agent compares two objects, she does so in accordance with the common preference.

Definition 2 (Common preference \succsim on A). For each $i \in N$ and $\{a, b\} \subseteq A_i$, $a R_i b \iff a \succsim b$.

Even if preferences are common in the sense that acceptable objects are compared in the same way, an object may be acceptable to some agents but not to others. Theorem 3 and its immediate corollary show that CONSTRAINEDWELFAREMAX remains NP-hard even with object-based weights and common preferences.

Theorem 3. *Deciding whether there exists a Pareto-efficient, welfare-maximizing allocation is NP-complete even for balanced problems with object-based weights and common preferences in which a single edge has non-zero weight.*

⁵ For Theorem 3, this follows by the same logic as for Theorem 1 as there again is only one non-zero weight. For Theorem 2, inapproximability can be shown through a construction similar to the one used in the proof of Theorem 2.

Table 1
Preferences over acceptable objects.

i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9
a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8	a_9
c_1	c_2	c_3	c_4	c_5	c_6			




Fig. 5. Acceptable pairs are connected by an edge.

Proof. We reduce from the NP-complete EXACT-3-COVER decision problem (Garey & Johnson, 1979). An instance of EXACT-3-COVER is as follows. We are given $C = \{c_1, \dots, c_{3n}\}$ and $B = \{B_1, \dots, B_m\}$ such that, for each $B_\ell \in B$, $B_\ell = \{c_1^\ell, c_2^\ell, c_3^\ell\} \subseteq C$. We wish to determine whether there is $B' \subseteq B$ such that each object in C is included in exactly one $B_\ell \in B'$. That is, is there a partition of C into n elements of B ? Next, we transform this instance of EXACT-3-COVER into its corresponding object allocation problem with preferences in which a single edge has non-zero weight. The size of the latter problem is not much larger than the original one (polynomial size in $m + n$). We show that the instance of EXACT-3-COVER has a feasible solution if and only if the non-zero weighted edge can be included in a Pareto-efficient allocation.

Define the total order \geq_C on C and, without loss, assume that each $B_\ell = \{c_1^\ell, c_2^\ell, c_3^\ell\}$ is labeled accordingly: $c_1^\ell >_C c_2^\ell >_C c_3^\ell$. For each $B_\ell = \{c_1^\ell, c_2^\ell, c_3^\ell\} \in B$, we create a subproblem (“gadget”) G_ℓ that includes objects B_ℓ and $A_\ell = \{a_1^\ell, \dots, a_9^\ell\}$ and agents $N_\ell = \{i_1^\ell, \dots, i_9^\ell\}$. For $k \neq \ell$, $N_k \cap N_\ell = A_k \cap A_\ell = A_\ell \cap C = \emptyset$. Moreover, no object in A_k is acceptable to an agent in N_ℓ . We label the objects in A_ℓ “gadget-specific”—they are only part of one gadget, and they are only acceptable to agents within this gadget. In contrast, the objects in C can be part of multiple gadgets and acceptable to any agent. We label them “common” objects. Preferences for an arbitrary gadget G_ℓ are in Table 1 (as all gadgets are symmetric, we drop the ℓ 's). Fig. 5 shows the associated acceptability graph.

Let $A_0 = \{a_9^1, \dots, a_9^m\}$ be the objects of “type” a_9^ℓ and define the total order \geq_{A_0} on A_0 . Add further a set N_0 of n agents who top-rank the objects in A_0 according to \geq_{A_0} followed by the objects in C according to \geq_C . Add also a special agent i^* and a special object a^* such that i^* top-ranks A_0 according to \geq_{A_0} , then C according to \geq_C , and last a^* . Hence, a^* is only acceptable to i^* . All in all, the object allocation problem contains agents $N = N_0 \cup N_1 \cup \dots \cup N_m \cup \{i^*\}$ and objects $A = C \cup A_1 \cup \dots \cup A_m \cup \{a^*\}$. (To ensure that the problem is balanced, add $2n$ agents who find no object acceptable.) Observe that there is a common preference \succsim on A : let \succsim rank the gadget-specific objects as $a_1 > \dots > a_9$ in order \geq_{A_0} and above the common objects, and then rank the common objects as in \geq_C , and last rank a^* . The object-based weights are such that the only non-zero weight is $w(i^*, a^*) > 0$. Hence, we wish to decide whether there exists a Pareto-efficient allocation that assigns a^* to i^* .

In order to assign a^* to i^* in a Pareto-efficient way, all objects in A_0 and C must also be assigned. Consider an arbitrary gadget G_ℓ . There is only one Pareto-efficient way of assigning a_9^ℓ to i_9^ℓ : namely, through $x_\ell = \{(i_1^\ell, a_1^\ell), \dots, (i_9^\ell, a_9^\ell)\}$. Alternatively, we can assign the common objects $c_1^\ell, c_2^\ell, c_3^\ell$ through y_ℓ :

$$y_\ell = \{(i_1^\ell, c_1^\ell), (i_2^\ell, c_2^\ell), (i_3^\ell, c_3^\ell), (i_4^\ell, a_1^\ell), \dots, (i_8^\ell, a_5^\ell), (i_9^\ell, a_7^\ell)\}.$$

We can then assign a_9^ℓ to an agent in N_0 .

Suppose that there exists a solution $B' \subseteq B$ to the instance of EXACT-3-COVER. Define its corresponding allocation x as follows. For each $B_\ell \in B'$, include the edges of y_ℓ in x together with an edge (j, a'_j) for some $j \in N_0$. For each $B_\ell \notin B'$, include x_ℓ in x . Finally, include the edge (i^*, a^*) in x . Then x is a constrained welfare-maximizing allocation. It covers the common objects C and n objects in A_0 in n gadgets, and then covers the remaining $m - n$ objects in A_0 in the remaining $m - n$ gadgets. In this way, we can assign a^* to i^* in a Pareto-efficient way.

In contrast, if there is no solution to the instance of EXACT-3-COVER, then we require more than n gadgets to cover C , leaving too few agents to cover A_0 . \square

Corollary 2. *CONSTRAINEDWELFAREMAX is NP-hard even for balanced problems with object-based weights and common preferences.*

Corollary 3 is immediate from the proof of Theorem 3 and complements (Saban & Sethuraman, 2015)) result on SD FEASIBILITY for the case of common but incomplete preferences.

Corollary 3. *SD FEASIBILITY is NP-complete even for balanced problems with common but incomplete preferences.*

3.2. Tractable cases

We proceed to conditions under which CONSTRAINEDWELFAREMAX is tractable. Theorem 4 shows that, in a sufficiently restricted setting, we can use a *Serial Dictatorship* to find a constrained welfare-maximizing allocation.

Theorem 4. *CONSTRAINEDWELFAREMAX is polynomial-time solvable for balanced problems with object-based weights and complete preferences.*

Proof. As the problem is balanced and the preferences complete, all objects are assigned at every Pareto-efficient allocation. As weights are object-based, every Pareto-efficient allocation creates the same welfare. The problem of finding a constrained welfare-maximizing allocation is then reduced to finding a Pareto-efficient allocation. This can be done efficiently through *Serial Dictatorship*. \square

A natural counterpart to common preferences is “common acceptability”: an object is acceptable to one agent whenever it is acceptable to all agents. However, an object that is not acceptable to anyone adds little to the problem, so we consider only the special case of complete preferences. When both common and complete, preferences are “the same” for all agents. Next, Theorem 5 shows that the Hungarian method can be used to find a constrained welfare-maximizing allocation when all agents have the same preference. Compared to Theorem 4, there is now a trade-off between Pareto-efficiency and welfare-maximization as the highest-valued objects may be the least liked.

Theorem 5. *CONSTRAINEDWELFAREMAX is polynomial-time solvable under common and complete preferences.*

Proof. Let $A^\circ \subseteq A$ be the maximal set of at most $|N|$ objects such that $a \succ b$ and $b \in A^\circ$ imply $a \in A^\circ$. Pareto-efficiency implies that all objects in A° must be assigned. To ensure that this is done, first adjust the weights by adding a large-enough constant $K > 0$ to $w(i, a)$ for each $i \in N$ and $a \in A^\circ$.

If $|A^\circ| < |N| < |A|$, then we need to assign some of the objects outside A° as well. Let $A^* \subseteq A$ be the minimal set of at least $|N|$ objects such that $a \succ b$ and $b \in A^*$ imply $a \in A^*$. In the acceptability graph, cut all edges to objects not in A^* . This ensures that only objects in A^* are assigned in the reduced problem. Apply the Hungarian method to find a matching of maximum weight in the reduced

graph. This is a constrained welfare-maximizing allocation in the original problem. \square

In addition, CONSTRAINEDWELFAREMAX is also tractable when all agents are indifferent between all objects acceptable to them. Such dichotomous preferences are a special case of common preferences.

Definition 3 (Dichotomous preferences). For each $i \in N$ and $\{a, b\} \subseteq A_i$, $a I_i b$.

Theorem 6. *CONSTRAINEDWELFAREMAX is polynomial-time solvable under dichotomous preferences.*

Proof. It is immediate that an allocation of maximum cardinality is Pareto-efficient under dichotomous preferences. To show the converse, suppose for contradiction that $x \in X$ is Pareto-efficient but that $y \in X$ is of larger cardinality than x . Then, in the symmetric difference between x and y , $(x \setminus y) \cup (y \setminus x)$, there is an alternating path that starts and ends with an edge in y . Updating x by assigning the objects along the path as in y allows one more agent to receive an object without another agent becoming unassigned. This is a Pareto-improvement, contradicting that x is Pareto-efficient. Next, add a constant $K > 0$ to all edge weights. For a sufficiently large K , the weights are such that an allocation of larger cardinality has higher welfare than a smaller allocation. Hence, a welfare-maximizing allocation in the new problem must be of maximum cardinality and therefore Pareto-efficient. In particular, the allocation is constrained welfare-maximizing in the original problem. \square

Furthermore, CONSTRAINEDWELFAREMAX is tractable when assigning a more preferred object leads to higher social welfare. Under this condition, the planner's interests are aligned with the agents' and the edge weights numerically represent the preferences.

Definition 4 (Aligned interests). For each $i \in N$ and $\{a, b\} \subseteq A_i$, $a R_i b \iff w(i, a) \geq w(i, b)$.

Theorem 7. *Under aligned interests, each welfare-maximizing allocation is Pareto-efficient.*

Proof. Suppose that $x \in X$ is welfare-maximizing but Pareto-dominated by $y \in X$. Hence, for each $i \in N$, $y_i R_i x_i$, and, for some $j \in N$, $y_j P_j x_j$. Under aligned interests, $y_i R_i x_i \iff w(i, y_i) \geq w(i, x_i)$ and $y_j P_j x_j \iff w(j, y_j) > w(j, x_j)$. But then $W(x) < W(y)$, contradicting that x is welfare-maximizing. \square

Hence, under aligned interests, finding a constrained welfare-maximizing allocation is reduced to finding a welfare-maximizing allocation. This can be done efficiently using the Hungarian method.

Given the weights w , let $M(w) \subseteq X$ be the set of welfare-maximizing allocations. There can be several such allocations from which we can make a particular selection through a small-enough perturbation of the edge weights. To do so, let $\Delta > 0$ be the welfare-difference between the welfare-maximizing $x \in M(w)$ and a second-best allocation:

$$\Delta = W(x) - \max_{y \in X \setminus M(w)} W(y).$$

Define the perturbed weights π such that $\pi(i, a) = w(i, a) + \delta(i, a)$, $\delta(i, a) \geq 0$, and $\sum_{(i,a) \in E} \delta(i, a) < \Delta$. It is immediate that $M(\pi) \subseteq M(w)$: the welfare of an allocation outside $M(w)$ has increased by less than Δ , so its welfare must remain smaller than that of those in $M(w)$. This intuition will be used to prove Theorem 8.

Theorem 8 shows that there exists a Pareto-efficient, welfare-maximizing allocation when the edge weights depend only on the

agent. Moreover, the allocation can be found efficiently by applying the Hungarian method to a particular perturbed problem. In contrast to [Theorem 7](#), welfare-maximizing allocations need now not be Pareto-efficient: for instance, permuting who gets what at a Pareto-efficient, welfare-maximizing allocation does not reduce welfare but may turn the allocation inefficient.

Definition 5 (Agent-based weights). For each $i \in N$ and $\{a, b\} \subseteq A_i$, $w(i, a) = w(i, b)$.

Theorem 8. Under agent-based weights, there exists a Pareto-efficient, welfare-maximizing allocation.

Proof. Create the perturbed π as follows from the agent-based w . Add a small amount to $w(i, a)$ for i 's least preferred object a ; then add a slightly larger amount to $w(i, b)$, i 's second least preferred object; and so on. Then $M(\pi) \subseteq M(w)$ and π satisfies aligned interests. By [Theorem 7](#), each $x \in M(\pi)$ is Pareto-efficient. \square

Finally, we derive an additional result under positive and object-based weights. [Theorem 9](#) shows that each constrained welfare-maximizing allocation is of maximum cardinality. That is, it assigns the maximum number of objects. Thus, though Pareto-efficiency may imply a loss in welfare, it can be obtained without a loss in terms of the size of the allocation.

Theorem 9. Under positive and object-based weights, each constrained welfare-maximizing allocation is of maximum cardinality.

Proof. Let $x \in X$ be a constrained welfare-maximizing allocation. As a first step, break ties "in favor of x ". That is, if agent i is indifferent between objects a and x_i , then replace i 's preference R_i by R'_i such that $x_i P'_i a$. Indifferences that do not pertain to x are broken arbitrarily. Replacing the preferences with ties R with the strict preferences R' (weakly) shrinks the set of Pareto-efficient allocation. However, as ties are broken in favor of x , x remains Pareto-efficient. Moreover, welfare is unchanged, so x remains constrained welfare-maximizing. Furthermore, the acceptability graph is unchanged, so the set of allocations and their cardinalities is unchanged.

To obtain a contradiction, assume that there is $y \in X$ of larger cardinality than x . Then, in the symmetric difference between x and y , $(x \setminus y) \cup (y \setminus x)$, there is an alternating path that starts and ends with an edge in y . Label the agents and objects of this path $i_1, a_1, \dots, i_n, a_n$. Without loss, assume that a_n is i_n 's most preferred object among those unassigned at x . Swap the objects along this path to create $z \in X$: specifically, set $z_{i_1} = a_1$, $z_{i_2} = a_2$, and so on, and otherwise $z_j = x_j$. As weights are object-based and positive, $W(z) = W(x) + w(i_n, a_n) > W(x)$. If z is Pareto-efficient, then this contradicts that x is constrained welfare-maximizing. Hence, z fails one of the following conditions ([Abraham et al., 2005](#); see also [Manlove, 2013](#), Section 6.2.2.1):

- Maximality (no unassigned object is acceptable to an unassigned agent);
- Trade-in-free (no agent prefers an unassigned object to her assigned object);
- Coalition-free (no group of agents can exchange their assigned objects in a Pareto-improving way).

There are fewer unassigned agents and objects at z than at x . Therefore, as x is Pareto-efficient and hence maximal, z is also maximal. If z is not coalition-free, then make it so by repeatedly performing Pareto-improving exchanges. As this does not change the set of assigned objects and the weights are object-based, welfare is unchanged. Assume that z is adjusted in this way until it is coalition-free.

Then, it only remains that z is not trade-in-free: some agent prefers an unassigned object to her assigned object. As x is Pareto-efficient, x is trade-in-free. Hence, agent $j \in N \setminus \{i_1, \dots, i_n\}$ cannot

prefer an unassigned object to her assigned object z_j , which either remains x_j or has improved further when making z coalition-free: thus, $z_j R_j x_j$ and there are fewer, in terms of set inclusion, unassigned objects (namely, a_n). Agent i_1 does not prefer an unassigned object at x when she is unassigned, hence they are not acceptable to her, so no unassigned object at z is acceptable to her. For i_n , $z_{i_n} = a_n$ is chosen specifically as her most preferred unassigned object at x (or improved further to make z coalition-free), so she cannot be upset. Hence, for some $1 < k < n$, i_k prefers some unassigned object b to her assigned object at z . Again, without loss, suppose that b is i_k 's most preferred unassigned object at x . Shorten the path to i_1, a_1, \dots, i_k, b and repeat the argument.

After a finite number of repetitions, the path no longer can be shortened. Neither the first nor the last agent of the path is upset, so if and when the path only consists of two agents, the corresponding $z \in X$ is Pareto-efficient. As noted, $W(z) > W(x)$, a contradiction. \square

4. Integer programming and competitive equilibrium

In this section, we provide a method for finding a constrained welfare-maximizing allocation in the general, unrestricted problem. Though this method is guaranteed to find a solution, it may, in the worst case, require an exhaustive search through the entire set of allocations.

To formulate the unrestricted problem as an integer program, we will need to describe Pareto-efficiency through a set of linear constraints. For this purpose, introduce a price $p_a \in \{0, \dots, |A|\}$ for each object a . Together with an allocation x , (x, p) is a **competitive equilibrium** ([Gale, 1960](#)) if the following conditions hold:

1. Each unassigned object a has price zero: $a \in A \setminus \cup_i x_i \Rightarrow p_a = 0$;
2. Each object b preferred to the assigned object a is more expensive: $b P_i a = x_i \Rightarrow p_a < p_b$;
3. Each object b equally good as the assigned object a is no cheaper: $b I_i a = x_i \Rightarrow p_a \leq p_b$.

[Theorem 10](#) is a variation on the well-known second welfare theorem. While similar conclusions have been established in related models (for instance, [Roth and Postlewaite, 1977](#), Theorem 1), we are not aware of any results that cover the case of [Theorem 10](#).

Theorem 10. Allocation $x \in X$ is Pareto-efficient if and only if there are prices $p \in \{0, \dots, |A|\}^A$ such that (x, p) is a competitive equilibrium.

Proof. Construct the directed envy-graph $(A, S \cup T)$ as follows (see [Aziz, Biró, Lang, Lesca, & Monnot, 2016](#); [Abraham et al., 2005](#); [Cechlárová, Eirinakis, Fleiner, Magos, Manlove, Mourtos, Oceláková, & Rastegari, 2016](#)). Each object $a \in A$ is a node. The two types of arcs (directed edges) S and T represent strict preference and ties, respectively. If the agent assigned object a prefers object b , so $b P_i a = x_i$, then include an arc $(a, b) \in S$ (that is, an arc from a to b). Similarly, if $b I_i a = x_i$, then include an arc $(a, b) \in T$. Furthermore, if both object a and agent i are unassigned, then include an arc $(a, b) \in S$ for each $b \in A_i$. Note that, if $a \in A_i$, then this creates a self-loop in the graph (x is not maximal). Finally, if object a is unassigned and there is $b P_i x_i$, then include an arc $(a, b) \in S$. Again, if $a = b$, then there is a self-loop (x is not trade-in-free). Then, the following conditions are equivalent:

- A. Allocation x is Pareto-efficient;
- B. There is no directed cycle in the envy-graph that contains at least one strict-preference arc;
- C. There are prices p such that (x, p) is a competitive equilibrium.

We proceed to show this equivalence in three steps.

$A \Rightarrow B$. This is immediate. If there is such a directed cycle, then swapping the objects along the cycle is a Pareto-improvement.

$B \Rightarrow C$. Remove the S -arcs and partition the nodes A into the strongly connected components A_1, \dots, A_m of the sub-graph (A, T) . Hence, objects a and b belong to the same class A_ℓ if and only if there is a path of T -arcs from a to b (and from b to a). There is no S -arc between two objects in the same class: if $\{a, b\} \subseteq A_\ell$ but $(a, b) \in S$, then there is a directed cycle with a strict-preference arc starting from (a, b) and continuing on the T -path from b back to a , a contradiction. Next, create a new directed graph $(A^*, (S^* \cup T^*))$ by replacing each strongly connected component A_ℓ of (A, T) by a single node $a_\ell \in A^* = \{a_1, \dots, a_m\}$ and include an arc $(a_k, a_\ell) \in S^*$ whenever there is $a \in A_k$ and $b \in A_\ell$, $a \neq b$, such that $(a, b) \in S$. Otherwise, if there is $a \in A_k$ and $b \in A_\ell$ such that $(a, b) \in T$, include an arc $(a_k, a_\ell) \in T^*$. Hence, connect two components if they contain objects connected in the envy-graph.

Suppose, to derive a contradiction, that $(A^*, S^* \cup T^*)$ has a directed cycle C . If C consists of T^* -arcs only, then the nodes of the corresponding components are strongly connected. Hence, they should belong to the same component; this is a contradiction to the design of $(A^*, S^* \cup T^*)$. If C instead includes an S^* -arc, then the envy-graph contains a forbidden cycle. Hence, there are no directed cycles in $(A^*, S^* \cup T^*)$. It is then well-known that there exists a topological order p on A^* such that $(a_k, a_\ell) \in S^* \Rightarrow p(a_k) < p(a_\ell)$ and $(a_k, a_\ell) \in T^* \Rightarrow p(a_k) \leq p(a_\ell)$. As there are at most $|A|$ nodes $a_\ell \in A^*$, we can fit each $p(a_\ell)$ in $\{0, \dots, |A|\}$. We extend this to a weak topological order on the original objects A by setting $p_a = p(a_\ell)$ for every $a \in A_\ell$.

Finally, we show that (x, p) is a competitive equilibrium. First, each unassigned object makes out its own component and has no incoming arcs. Therefore, its price can be set to zero. Second, if agent i prefers object b to $x_i = a$, so $(a, b) \in S$, then $p_a < p_b$. Third, if agent i is indifferent between b and $x_i = a$, then $p_a \leq p_b$, as required.

$C \Rightarrow A$. To derive a contradiction, suppose that x is not Pareto-efficient. First, suppose that there is an improving path that ends with an unassigned object. This object must have price zero. Furthermore, the agent i who wants to exchange x_i for the unassigned object must be indifferent between them: otherwise, the price of x_i is negative. This continues along the cycle, all the way to the first agent: at each step, the agent is indifferent between the object that she is assigned at x and the object that she “points to” in the cycle. Then no agent is better off (in the strict sense), a contradiction to it being an improving path.

Suppose instead that there is a Pareto-improving exchange. As we move through the cycle, the objects’ prices must be non-decreasing: this as each agent “points” to an object at least as good as the one they are assigned at x . But as we move through the cycle, we eventually end up where we started, so prices cannot increase either. Hence, all prices are equal. But then, again, no agent is better off in the strict sense. \square

Theorem 10 allows us to reformulate the problem of constrained welfare-maximization: we want to find the allocation of highest welfare among those that can be supported by some prices in a competitive equilibrium. We proceed to formulate the corresponding integer program. Let the binary decision variables $\mathbf{x} \in \{0, 1\}^E$ indicate assignment: for each edge $(i, a) \in E$, $x_{ia} \in \{0, 1\}$ is such that agent i is assigned object a whenever $x_{ia} = 1$. Let the integer decision variables $p \in \{0, \dots, |A|\}^A$ denote prices. The objective is to maximize welfare:

$$W(\mathbf{x}) = \sum_{(i,a) \in E} x_{ia} w(i, a). \tag{1}$$

We introduce some additional variables to help formulate the constraints. Let the binary decision variables c indicate whether an agent/object is covered in the allocation. For agent i and object a ,

$$c_i = \sum_{a \in A_i} x_{ia} \quad \text{and} \quad c_a = \sum_{i \in N} x_{ia}. \tag{2}$$

Let the binary variables s and t be derived from the preference data. In particular, say agent i is assigned object a but finds object b at least as good. If this is in the strict sense, $b P_i a$, then $s_{ab} = 1$; if this is through a tie, $b I_i a$, then $t_{ab} = 1$:

$$s_{ab} = \sum_{\substack{i \in N: \\ b P_i a}} x_{ia} \quad \text{and} \quad t_{ab} = \sum_{\substack{i \in N: \\ b I_i a}} x_{ia}. \tag{3}$$

To complete the formulation, we introduce the constraints. No agent $i \in N$ receives more than one object and no object $a \in A$ is assigned to more than one agent:

$$c_i \leq 1 \quad \text{and} \quad c_a \leq 1. \tag{4}$$

A Pareto-efficient allocation is necessarily maximal. Therefore, for each edge $(i, a) \in E$, at least one of i and a is covered in the allocation:

$$c_i + c_a \geq 1. \tag{5}$$

The remaining constraints ensure that (\mathbf{x}, p) is a competitive equilibrium. First, object $a \in A$ has zero price if it is unassigned:

$$c_a |A| \geq p_a. \tag{6}$$

If $c_a = 0$, then $p_a = 0$ as $p_a \geq 0$; if $c_a = 1$, then the constraint is always satisfied as $p_a \leq |A|$. Second, if the agent assigned a prefers b , then b should be more expensive than a . For each $\{a, b\} \subseteq A$,

$$(1 - s_{ab})(|A| + 1) + p_b \geq p_a + 1. \tag{7}$$

If $s_{ab} = 1$, then $p_b \geq p_a + 1$, so $p_b > p_a$; if $s_{ab} = 0$, then the constraint is always satisfied as $|A| \geq p_a$ and $p_b \geq 0$. Third, if the agent instead is indifferent between a and b , then b should be no cheaper than a :

$$(1 - t_{ab})|A| + p_b \geq p_a. \tag{8}$$

This completes our formulation. **Theorem 11** follows immediately.

Theorem 11. Let $(\mathbf{x}^*, p^*) \in \{0, 1\}^E \times \{0, \dots, |A|\}^A$ maximize (1) subject to (2) through (8). Then $\{(i, a) \in E \mid x_{ia}^* = 1\}$ is a constrained welfare-maximizing allocation.

We can extend the integer program to incorporate that some object is available in multiple copies, say to capture quotas in school choice. If we instead treat (some of) the seats symmetrically, then we can modify the program as follows.

Let $q_a \in \mathbb{N}$ denote the quota of object a . The previously binary variables c , s , and t will now be integer-valued and we adjust the corresponding constraints as follows. Eq. (4) is replaced by $c_a \leq q_a$. Furthermore, we introduce new indicator variables f_a , \tilde{s}_{ab} , and \tilde{t}_{ab} to replace c_a , s_{ab} , and t_{ab} in the formulas in the following way. We introduce $f_a \in \{0, 1\}$ to indicate whether a is fully assigned: $c_a < q_a \iff f_a = 0$. This is achieved through the following constraints:

$$f_a q_a \leq c_a \quad \text{and} \quad f_a + q_a \geq c_a + 1.$$

Eqs. (5) and (6) are replaced by $c_i + f_a \geq 1$ and $f_a |A| \geq p_a$, respectively. We also introduce $\tilde{s}_{ab} \in \{0, 1\}$ to indicate whether s_{ab} is positive. To do so, add the constraints $\tilde{s}_{ab} \leq s_{ab}$ and $\tilde{s}_{ab} |N| \geq s_{ab}$. Replace s_{ab} by \tilde{s}_{ab} in Eq. (7). Finally, make the analogous change for t_{ab} : introduce its indicator \tilde{t}_{ab} together with the two associated constraints, and replace t_{ab} by \tilde{t}_{ab} in Eq. (8).

Table 2
Results for the 2016 kindergarten allocation in Harku, Estonia.

Rank	Stable		Pareto-efficient									
	DA		IA		TTC		CWM		OPDA		WM	
	Pref.	Dist.	Pref.	Dist.	Pref.	Dist.	Pref.	Dist.	Pref.	Dist.	Pref.	Dist.
1	102	103	123	95	102	98	106	102	98	104	80	102
2	23	26	6	29	20	24	23	32	22	24	30	29
3	5	3	1	3	6	5	4	3	9	6	14	8
4	8	7	7	7	11	7	4	6	9	6	10	6
5	5	4	9	11	4	4	6	6	7	5	6	7
6	7	2	5	3	6	3	8	1	6	2	8	0
7	2	7	1	4	3	11	1	2	1	5	4	0
Average rank	1.82	1.8	1.63	1.91	1.85	2	1.74	1.64	1.86	1.75	2.16	1.6
Average distance	3398		3810		3789		3171		3319		3024	
Blocking agents	0		21		34		26		31		45	
Swaps in post-TTC	2		0		0		0		14		37	

5. Estonian kindergarten allocation

In this section, we examine the 2016 kindergarten allocation in Harku, Estonia (see Veski et al., 2017) in which each of 152 children were to be assigned one out of 155 seats across seven kindergartens. The data contains the families’ stated preferences and the travel distances to the various kindergartens. In roughly 81% of the cases, a closer school is preferred (hence, aligned interest is at least partly satisfied). The edge weight between family i and kindergarten a is $w(i, a) = D - d(i, a)$, where D is the maximum distance in the data and $d(i, a)$ is the distance between i and a . This ensures that edge weights are non-negative and higher for students living closer. We examine the stable *Deferred Acceptance* (DA; the student- and school-proposing versions yield the same allocation) and the Pareto-efficient *Immediate Acceptance* (IA) and *Top Trading Cycles* (TTC). We give priority to families living closer to the kindergarten and break ties randomly. We contrast these solutions with the constrained welfare-maximizing solution (CWM), the unrestricted welfare-maximizing solution (WM), and *Optimal Priority Deferred Acceptance* (OPDA, see Appendix A). Table 2 shows the results which we discuss next (see Diebold and Bichler, 2017, for a similar comparison of algorithms).

For each solution, we determine the number of families assigned their top-ranked kindergarten, their second-highest ranked kindergarten, and so on. Similarly, we count the number of families assigned their closest kindergarten, second closest, and so on. IA stands out by assigning most families their top choice but also fewest families their closest kindergarten. OPDA, DA, CWM, and WM assign most families to close kindergartens. This is confirmed by the average ranks for which IA and CWM/WM stand out with regard to the preferences and distances, respectively. By construction, WM minimizes the average distance. Imposing Pareto-efficiency (CWM) increases the average distance by roughly 150 meters, but CWM still performs considerably better than the priority-based solutions. The final two rows capture instability and inefficiency. CWM reduces the number of blocking families compared to TTC and WM. At CWM, a majority of blocking families prefer a more distant kindergarten to the one at which they are assigned. Moreover, there is an unpopular kindergarten that never fills its seats and many of the children placed there have justified envy. For the final row, we first determine the respective solutions and then apply TTC to find Pareto-improving swaps or cycles (so the final allocation is Pareto-efficient). DA only requires two families to swap assignments for the allocation to be Pareto-efficient, whereas OPDA and WM require much greater changes.

The case study also shows that CWM is operational in practice: even though the IP used to solve CWM contains more than 1,000 variables and constraints, the solution is obtained within seconds. Simulated problems many times the size of the case study are also solvable in reasonable time (see Appendix B); turning to commercial solvers and using more sophisticated techniques (Firat et al., 2016) would likely push the boundary yet further.

While the improvement that CWM provides in terms of social welfare is important, as shown by the court case of Example 1 and the Boston school choice redesign to reduce busing costs, the solution may also have some drawbacks. Specifically, it is well-known that solutions such as DA, TTC, and OPDA incentivize truthful reporting of preferences (Abdulkadiroğlu and Sönmez, 2003) while IA is manipulable. Next, we examine the strategic properties of CWM and address whether it is realistic to expect agents to state their true preferences when the solution is used to allocate the objects.

6. Incentives

An agent’s preference is typically private information that she reveals to the planner. In this section, we examine whether we can select constrained welfare-maximizing allocations in such a way that no agent ever benefits from misstating her preference. We first find a positive result when the problem is sufficiently restricted, and then show that manipulation is possible if these restrictions are relaxed. For now, we take as given that the weights are set independently of the reported preferences.

To derive a positive result, we refer back to Theorem 4: for balanced problems with object-based weights and complete preferences, *Serial Dictatorship* can be used to efficiently find a desirable allocation. As *Serial Dictatorship* is not manipulable, we can select constrained welfare-maximizing allocations in a non-manipulable way under these restrictions. Note that this considers only manipulations achieved by shuffling the preference list (that is, all objects are always acceptable).

Next, we relax each of the three conditions, one at a time. First, Example 3 shows a beneficial manipulation when there are more objects than there are agents.

Example 3. Let $N = \{1, 2\}$, $A = \{a, b, c\}$, and $R_1 = R_2$ be such that $a P_1 b P_1 c$. Edge weights are object-based with $w(i, b) > w(i, c)$. Then $x = (a, b)$ and $y = (b, a)$ are constrained welfare-maximizing. Without loss, suppose that x is selected. Then agent 2 benefits from reporting R'_2 such that $a P'_2 c P'_2 b$: the unique constrained welfare-maximizing allocation at (R_1, R'_2) is y .

Second, we relax complete preferences. For the positive result, we considered only manipulations achieved by shuffling the preference list. That is, agents could not alter whether an object was acceptable or not. [Example 4](#) shows that the result is overturned when agents can state that some object is unacceptable.

Example 4. Let $N = \{1, 2\}$, $A = \{a, b\}$, and $R_1 = R_2$ be such that $a P_1 b$. Edge weights are object-based. Then $x = (a, b)$ and $y = (b, a)$ are constrained welfare-maximizing. Without loss, suppose that x is selected. Then agent 2 benefits from stating that b is unacceptable: the unique constrained welfare-maximizing allocation is then y .

Third, [Example 5](#) shows that an agent may manipulate if weights no longer are object-based.

Example 5. Let $N = \{1, 2, 3\}$, $A = \{a, b, c\}$, and $R_1 = R_2 = R_3$ be such that $a P_1 b P_1 c$. The unique non-zero edge weight is $w(3, c) = 1$. Then $x = (a, b, c)$ and $y = (b, a, c)$ are constrained welfare-maximizing. Without loss, suppose that x is selected. Then agent 2 benefits from reporting R'_2 such that $a P'_2 c P'_2 b$: the unique constrained welfare-maximizing allocation at (R_1, R'_2, R_3) is y .

These findings are summarized in [Theorem 12](#).

Theorem 12. *Under object-based weights, complete preferences, and with at least as many agents as objects, a constrained welfare-maximizing allocation can be selected in a non-manipulable way. Relaxing either of these conditions may allow beneficial manipulation.*

Whereas we so far have assumed that the weights are set independent of the reported preferences, we would actually encourage the practitioner to do otherwise (see also [Chiarandini, Fagerberg, and Gualandi, 2017](#), for a comprehensive study of weighting preferences). For instance, in practice we may be limited to allocations which include every agent. Such a constraint may clearly distort incentives as a simple option then is to report only the agent's most preferred object. This strategy may backfire if too many agents adopt it, but otherwise it is an easy way to ensure one's top choice. A more interesting approach is to encourage agents to report complete preferences, which would make it easier to find everyone an object. The designer may announce that, whenever an agent reports a complete preference, then the agent's edge weights are increased. Not only does this give incentives to report complete preferences, but it also reduces the possibilities of agents manipulating by truncating their preferences. Formally, once the weights depend on the stated preferences, the preference revelation game is very different. Depending on how it is set up, it may now either be easier or even impossible to manipulate.⁶ For large problems in which it is impractical to report complete preferences, a natural alternative is to require agents to report preferences of (at least) a particular length. A possible manipulation may then be to fill the reported preference with objects the agent is unlikely to get: the agent top-ranks her favorite object and then lists only very popular objects or objects bad with respect to the objective of the planner (such as far-away schools).

Even if a mechanism is manipulable in theory, it does not necessarily mean that it will be manipulated in practice. And conversely, laboratory experiments have shown that, even for a provably non-manipulable solution such as DA, experimental subjects do not always report preferences truthfully ([Chen & Sönmez, 2006](#)). There are some new theoretical concepts for addressing this behavior, such as strategy-proofness in the large

([Azevedo & Budish, 2019](#)) and obvious strategy-proofness ([Li, 2017](#)). Moreover, in line with the conclusions of [Chiarandini et al. \(2017\)](#), the NP-hardness of the problem and the random selection among constrained welfare-maximizing solutions may make it harder for agents to manipulate. To get a better understanding of the manipulability of our solution, a more thorough Bayesian approach may be more realistic (but also more challenging) to use, where we instead would consider the expected gains of manipulation.

How would one manipulate the constrained welfare-maximizing solution? In some cases, manipulation is counter-intuitive (an agent may manipulate by reversing her preference), but we can find a simple strategy by returning to the Estonian kindergarten allocation of [Section 5](#). In this case study, there is one unpopular kindergarten that fails to fill its seats. By Pareto-efficiency, no child is therefore assigned somewhere less preferred than the unpopular kindergarten. Therefore, a child living relatively far from the unpopular kindergarten and relatively close to her preferred kindergarten is often able to guarantee herself her first choice by ranking the unpopular kindergarten second. Among the 46 families not assigned their most preferred choice, a majority are able to get their first choice by following this strategy (we manipulate for one family at a time, keeping all other families' preferences unchanged) but there are also some for whom it does not pay off and the child instead is placed at the unpopular kindergarten.

In summary, for understanding the actual manipulability of CWM, one would need to further study its properties through sophisticated theoretical models and also test its practical manipulability through laboratory and field experiments.

7. Concluding remarks

We have studied the problem of finding a constrained welfare-maximizing allocation, that is, a Pareto-efficient allocation of highest welfare. This problem is NP-hard even under strong conditions, but we have also identified settings in which the problem is polynomial-time solvable. For the general problem, we formulated an integer program. We used this program to solve a real-world instance of kindergarten admissions in which edge weights represent travel distance, and it was quick to determine the solution. However, there are still many open questions left to study.

Implicit throughout has been that Pareto-efficiency takes precedence over welfare. (Under aligned interests or agent-based weights, the order is irrelevant: we can attain Pareto-efficiency and welfare-maximization simultaneously.) That is, the planner wants to select the Pareto-efficient allocation of highest welfare. A different approach is to define a measure of how Pareto-efficient an allocation is and then select a welfare-maximizing allocation of "highest Pareto-efficiency". This is left for future research.

Also outside the scope of the current paper, an interesting extension is to equip objects with "standard" priorities over agents and, say, find the stable allocation of highest welfare. Extending the model is unlikely to overturn the negative results but can be interesting for the positive results. For instance, an extension of the condition imposed in [Theorem 7](#), aligned interests, is to align the added priorities with the edge weights.

Another extension is to allow agents to receive multiple objects. In this case, it is already challenging to elicit the preferences of the agents over the bundles, as the number of bundles is exponential in the number of objects. This is often resolved by eliciting ordinal or cardinal preferences over single objects, and then extending this to preferences over bundles through, for instance, responsiveness or additivity. However, in the cardinal setting it can already be computationally hard to decide whether an allocation is Pareto-efficient ([Aziz et al., 2016](#)). But in line with our positive results in

⁶ Taken to its extreme, once preferences influence weights, any solution can be obtained as welfare-maximizing. For instance, to select the outcome of DA, then we can set the weights to 0 or 1 depending on whether the agent is assigned the object under DA. However, for all intents and purposes, this completely removes the intended interpretation of "welfare maximization".

Section 3, finding constrained welfare-maximizing solutions may be still tractable under some conditions.

Appendix A. Solutions: Definitions and algorithms

In this section, we provide more formal definitions of the solutions referred to throughout the paper. For the reader interested in learning more, we refer to the excellent surveys by Manlove (2013) and Haeringer (2017).

We adopt the terminology of Section 5 and refer to agents as students and objects as seats at schools. We impose some restrictions compared to the model introduced in Section 2. First, every school is acceptable to every student. Second, students have strict preferences over schools (no indifferences). Third, each school a has q_a seats (its “quota”) and there are sufficiently many seats in total to assign every student. Finally, weights represent the distances between the students and the schools as in Section 5.

Most of the mechanisms are *priority-based* in the sense that students have different priorities at different schools. We examine two types of priorities. First, “distance-based priorities” give higher priority to students living closer to schools. Second, “distance-based priorities adjusted by allocation x ” prioritize students assigned to the school under x to those not and otherwise prioritize on distance. As an example, suppose that students i , j , and k live 100, 200, and 1000 meters from school a . Distance-based priorities yield the order i, j, k with student i given the highest priority. Suppose further that the allocation x assigns both students j and k , but not i , to school a . Adjusting the distance-based priorities to x then results in the order j, k, i . In what follows, priorities are adjusted for two reasons. First, by adjusting priorities to x and then executing *Top Trading Cycles*, we obtain a measure of how far from Pareto-efficient x is. Second, by adjusting priorities to the welfare-maximizing allocation, we can study solutions that “lie between” the original solution (say, *Deferred Acceptance*) and the welfare-maximizing solution. This will be used to define the new *Optimal Priority Deferred Acceptance*.

Top Trading Cycles. This mechanism is defined using a directed bipartite graph. Nodes are given by the students and schools. Each student has an outgoing arc to her most preferred school; each school has an outgoing arc to the student with highest priority at the school. The graph contains at least one cycle. Each student in the cycle is assigned to the school she points to and then removed from the graph. Similarly, school a is removed from the graph once it has filled its quota q_a . In succeeding rounds, students previously pointing to school a point to their most preferred school among those that remain in the graph. Schools redirect their arcs similarly to the remaining student with highest priority. If there are several cycles at some stage, the cycles not selected remain cycles in the subsequent round. Therefore, the outcome is independent of the order in which the cycles are selected.

In Section 5 and Appendix B, we report values on “swaps in post-TTC”. These are computed as follows. We adjust priorities (as described above) to whichever allocation that we are examining, and we then execute *Top Trading Cycles*. Whenever we process a cycle of at least two students, we keep track of how many students change schools. The number of “swaps in post-TTC” is the total number of changes across all cycles.

Immediate Acceptance. In this mechanism, students “propose” to their most preferred schools. If a school can seat all its proposers, then it does so. A school receiving more proposal than it has seats “immediately accepts” the proposals from the students with highest priority and rejects the others. All rejected students proceed to propose to their most preferred school that still has vacant seats (“Immediate Acceptance with Skips” in Harless, 2019).

Deferred Acceptance. This operates as *Immediate Acceptance* with the exception that the allocation is not made final until at the

a ————— 1 — b ————— 2

Fig. A.6. Students and schools in Example 6.

end of the algorithm. Students again propose to their most preferred schools. Each school then tentatively accepts proposers up to its quota. Rejected students proceed to propose to their most preferred school that has not yet rejected them. Each school then, again, tentatively accepts proposers up to its quota—choosing both from those newly proposing and those tentatively accepted in the previous round. In this way, a school may tentatively accept a student at first only to reject the student at a later stage.

Optimal priority deferred acceptance.⁷ This mechanism is implemented in two steps. First, we compute the welfare-maximizing allocation and adjust the priorities on the basis of it. That is to say, students assigned school a under the welfare-maximizing allocation are prioritized over those not; otherwise, students are prioritized on distance. From a practical perspective, this step can be undertaken as soon as the geographical distribution of the students is known (for larger cities, the distribution likely only changes slowly over time). The allocation can then be announced to the students in the form of “catchment areas” (compare Fig. B.7 in Appendix B). That is, we announce to the students that there has been an initial assignment according to the catchment areas. Students are guaranteed seats at their respective schools, but if they prefer to relocate elsewhere, they are given the option to submit preferences over preferred schools. At that point, *Deferred Acceptance* is executed and identifies Pareto-improvements over the initial assignment. Put succinctly, *Optimal Priority Deferred Acceptance* is *Deferred Acceptance* executed on the instance with priorities adjusted to the welfare-maximizing allocation.

The mechanism results in an allocation that is stable with respect to the adjusted priorities. Recall that a stable allocation is one that refuses a student a preferred school only if all its seats are assigned to higher-priority students. In this case, if student i prefers school a over her assignment, then a is filled with students who all either are from a 's catchment area or who live closer to a than i does. It is immediate that *Optimal Priority Deferred Acceptance*, like *Deferred Acceptance*, cannot be manipulated. Moreover, as both finding the welfare-maximizing allocation and executing *Deferred Acceptance* can be done in polynomial time, we quickly obtain the solution. Among its drawbacks, the solution is neither Pareto-efficient nor stable with respect to the original distance-based priorities.

Alternatively, the solution can be implemented by first having the students report their preferences, then adjusting the priorities for the welfare-maximizing allocation, and finally executing *Deferred Acceptance*. In this case, it is important that the preferences do not affect the welfare-maximizing allocation. Specifically, even if student i reports school a as unacceptable, i should still be allowed to be assigned to a when the welfare-maximizing allocation is determined. If not, if we look for the welfare-maximizing allocation that assigns students to acceptable schools only, then students will be able to manipulate the solution. This is illustrated in Example 6, which also shows that *Optimal Priority Deferred Acceptance* may unnecessarily leave students unassigned. To summarize, provided that the welfare-maximizing allocation is computed independently of the reported preferences, it does not matter for the solution and its strategic properties whether we collect preferences before or after we determine priorities. However, the benefit of first determining priorities is that it assigns each student a guaranteed school. The students then only need to report schools preferred to the guaranteed school. In the Estonian

⁷ We are grateful to Philippe Jehiel for suggesting this mechanism to us.




Fig. B.7. Illustration of a simulated instance with 560 students and eight schools, each with 70 seats. Large markers are schools, small markers of the same color are students assigned to the different schools. In the left figure, the inner circle is the student's assigned school under the welfare-maximizing solution; the outer circle is the assigned school under OPDA. In the right figure, the inner circle is the student's most preferred school while the outer circle again is the student's assignment under OPDA.

case, 80 out of 152 families need not report anything—they are assigned their preferred school under WM. Out of the remaining 72 families, 30 only have to report their top choice. In contrast, if we collect preferences first, then all students may have to report complete preferences.

Example 6. Consider students 1 and 2 together with schools a and b arranged along a line as in Fig. A.6. The parents of 1 work “to the right” and only accept dropping of 1 at school b . Hence, a is unacceptable. The parents of 2 work “to the left” and prefer b slightly over a . Still, a remains acceptable. We assume further that 1 lives closer to b than 2 does.

Consider first the welfare-maximizing solution that ignores the acceptability constraints. It selects the allocation $\{(1, a), (2, b)\}$. When we adjust priorities accordingly and execute *Deferred Acceptance*, student 2 prefers and has highest priority at b . Hence, we assign 2 to b . As *Deferred Acceptance* does not assign students to unacceptable schools, student 1 remains unassigned. Hence, *Optimal Priority Deferred Acceptance* selects $\{(2, b)\}$.

Consider next the welfare-maximizing solution that only assigns students to acceptable schools. Assuming this is set up to include a large cost to leaving a student unassigned, it selects the allocation $\{(1, b), (2, a)\}$, which also is the output of the following execution of *Deferred Acceptance*. Hence, *Optimal Priority Deferred Acceptance* selects $\{(1, b), (2, a)\}$.

Consider finally the case when student 1 prefers school b but is willing to accept both schools. The welfare-maximizing allocation will then be as in the first case, $\{(1, a), (2, b)\}$, which also will be the final output of *Optimal Priority Deferred Acceptance*. In this way, if the first step of *Optimal Priority Deferred Acceptance* is set up to only assign students to acceptable schools, then student 1 can manipulate by reporting that a is unacceptable.

Appendix B. Simulations

We run a simulation study to further compare the different solutions.⁸ Each instance is composed of 1000 students and 10 schools with 100 seats each. Every student is acceptable to every school, so all solutions output complete allocations in which no student is left unassigned. Geographical coordinates for the students and schools are drawn randomly within the unit circle through a distance (uniformly from $[0,1]$) and angle (uni-

formly from 0 to 360 degrees) from the center. School priorities are distance-based, giving higher priority to student living closer (measured by the Euclidean distance). Weights are set up as in Section 5: the edge weight between student i and school a is $w(i, a) = D - d(i, a)$, where D is the maximum distance in the data and $d(i, a)$ is the distance between i and a . Finally, each school is assigned a quality level uniformly from $[0,1]$.

Student preferences are derived through a linear combination of school quality κ , distance d , and a random noise term ε . The utility that student i derives from being assigned to school a is positively affected by the quality of a and negatively by the distance:

$$u(i, a) = \alpha_\kappa \cdot \kappa(a) - \alpha_d \cdot d(i, a) + \alpha_n \cdot \varepsilon(i, a).$$

The parameters α_κ , α_d , and α_n control how much weight is put on quality, distance, and the random noise, respectively. As an example, with $\alpha_d = 1$ and $\alpha_\kappa = \alpha_n = 0$, preferences are completely distance-based and we obtain the case of *aligned interests*. If instead $\alpha_\kappa = 1$ and $\alpha_d = \alpha_n = 0$, we obtain the case of *common preferences*. Finally, with $\alpha_\kappa = \alpha_d = 0$ and $\alpha_n = 1$, preferences are uncorrelated with distance and school quality.

In Table B3, we report the simulation results for three different parameters settings, in each case averaged across 100 instances. The top rows labelled “distance” refer to primarily distance-based preferences with parameters $\alpha_d = 3/5$ and $\alpha_\kappa = \alpha_n = 1/5$. The middle rows, “quality”, refer to primarily quality-based preferences with parameters $\alpha_d = \alpha_n = 1/5$ and $\alpha_\kappa = 3/5$. For the final rows, “random”, we set $\alpha_d = \alpha_\kappa = \alpha_n = 1/3$. Among the results in Table B3, we wish to highlight the following:

1. The average distance is very similar for CWM and OPDA under all three parameter settings and always considerably better than for DA, IA, and TTC. This is in some contrast to Section 5, in which OPDA and DA produced similar results while CWM significantly reduced the average distance.
2. The average preference ranks are considerably worse than in Section 5. That is, in the Estonian case study, more students are assigned their top schools. The fraction of blocking agents is considerably higher here than in Section 5. To explain for instance the high number for TTC under quality-based preferences, imagine that student i living close to popular school a wishes to go elsewhere—possibly to the distant school b . In the cycle that assigns i to b , a student j living close to b may get assigned to the popular school a . Due to the limited number of seats at a , a lot of students living closer to a than i does will not be assigned to a and therefore block the allocation.

⁸ The simulations are run in Python 3 with the free open source software PuLP using its default CBC solver.

Table B3
Results for the simulations.

		Stable	Pareto-efficient				
		DA	IA	TTC	CWM	OPDA	WM
Distance	Average preference rank	2.16	2.09	2.15	1.93	2.05	2.12
	Average distance rank	2.27	2.31	2.34	1.76	1.74	1.7
	Average distance	36,758	36,922	37,351	31,580	31,443	31,203
	Blocking agents	0	154	195	402	434	460
	Swaps in post-TTC	11	0	0	0	112	158
Quality	Average preference rank	4.17	3.89	4.08	4.1	4.27	4.37
	Average distance rank	2.5	2.82	2.83	1.84	1.81	1.74
	Average distance	39,072	42,274	42,331	32,565	32,375	31,745
	Blocking agents	0	422	606	642	639	649
	Swaps in post-TTC	73	0	0	0	166	213
Random	Average preference rank	2.76	2.56	2.69	2.68	2.96	3.38
	Average distance rank	2.59	2.84	3.03	2	1.98	1.77
	Average distance	39,772	42,283	44,272	33,788	33,549	31,904
	Blocking agents	0	246	488	535	554	609
	Swaps in post-TTC	80	0	0	0	239	372

- As in Section 5, the average preference rank is often lowest for IA. Among the other solutions, CWM typically assigns students to more preferred schools than OPDA does. Still, OPDA is an improvement over WM.
- The average distance rank is almost identical for CWM and OPDA.
- In terms of number of Pareto-improving swaps, OPDA is an improvement over WM but considerably worse than DA.

Figure B.7 graphically illustrates the relation between WM, OPDA, and the students' preferred schools for a simulated instance. In the left figure, comparing WM and OPDA, we can note the Pareto-improving swaps between the blue schools as well as the yellow and dark schools. These students live closer to the school they are assigned under WM but prefer the more distant school that they are assigned under OPDA. In the right figure, we see that the pink school is of high quality. Therefore, students who live close to it also are likely to prefer it. Confirming this intuition, WM and OPDA coincide for the pink school in the left figure.

References

- Abdulkadiroğlu, A., Che, Y. K., & Yasuda, Y. (2011). Resolving conflicting preferences in school choice: The "Boston Mechanism" reconsidered. *American Economic Review*, 101(1), 399–410.
- Abdulkadiroğlu, A., Che, Y. K., & Yasuda, Y. (2015). Expanding "choice" in school choice. *American Economic Journal: Microeconomics*, 7(1), 1–42.
- Abdulkadiroğlu, A., Che, Y. K., Pathak, P. A., Roth, A. E., Tercieux, O. (2017). Minimizing justified envy in school choice: The design of New Orleans' OneApp. No. w23265, National Bureau of Economic Research.
- Abdulkadiroğlu, A., Pathak, P. A., & Roth, A. E. (2005a). The New York City High School match. *American Economic Review*, 95(2), 364–367.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., & Sönmez, T. (2005b). The Boston Public School match. *American Economic Review*, 95(2), 368–371.
- Abdulkadiroğlu, A., & Sönmez, T. (1998). Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3), 689–701.
- Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, 93(3), 729–747.
- Abraham, D. J., Cechlárová, K., Manlove, D. F., & Mehlhorn, K. (2005). Pareto optimality in house allocation problems. In *Proceedings of the international symposium on algorithms and computation* (p. 1163–1175). Springer.
- Ágoston, K. C., Biró, P., & McBride, I. (2016). Integer programming methods for special college admissions problems. *Journal of Combinatorial Optimization*, 32(4), 1371–1399.
- Alpern, S., & Katrantzi, I. (2009). Equilibria of two-sided matching games with common preferences. *European Journal of Operational Research*, 196(3), 1214–1222.
- Andersson, T. (2017). Matching practices in elementary schools – Sweden, MiP country profile 24. <http://www.matching-in-practice.eu/elementary-schools-in-sweden/>.
- Andersson, T., & Ehlers, L. (2017). Assigning refugees to landlords in Sweden: Efficient stable maximum matchings. *The Scandinavian Journal of Economics*. Wiley Online Library.
- Ashlagi, I., & Shi, P. (2016). Optimal allocation without money: An engineering approach. *Management Science*, 62(4), 1078–1097.
- Aygün, O., & Bó, I. (2017). *College admission with multidimensional privileges: The Brazilian affirmative action case*. Available at SSRN 3071751.
- Aygün, O., & Turhan, B. (2017). Large-scale affirmative action in school choice: Admissions to IITs in India. *American Economic Review Papers and Proceedings*, 107(5), 210–213.
- Azevedo, E. M., & Budish, E. (2019). Strategy-proofness in the large. *Review of Economic Studies*, 86(1), 81–116.
- Aziz, H., Biró, P., Lang, J., Lesca, J., & Monnot, J. (2016). Optimal reallocation under additive and ordinal preferences. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems* (pp. 402–410). International Foundation for Autonomous Agents and Multiagent Systems.
- Bansak, K., Ferwerda, J., Hainmueller, J., Dillon, A., Hangartner, D., Lawrence, D., & Weinstein, J. (2018). Improving refugee integration through data-driven algorithmic assignment. *Science*, 359(6373), 325–329.
- Bansak, K., Hainmueller, J., & Hangartner, D. (2016). How economic, humanitarian, and religious concerns shape European attitudes toward asylum seekers. *Science*, 354(6309), 217–222.
- Bansak, K., Hainmueller, J., & Hangartner, D. (2017). Europeans support a proportional allocation of asylum seekers. *Nature Human Behaviour*, 1(0133).
- Biró, P. (2011). University admission practices – Hungary, MiP country profile 5. <http://www.matching-in-practice.eu/higher-education-in-hungary/>.
- Biró, P. (2017). Applications of matching models under preferences. In U. Endriss (Ed.), *Trends in Computational Social Choice*. AI Access, 18, 345–373.
- Biró, P., Fleiner, T., & Irving, R. W. (2016). Matching couples with Scarf's algorithm. *Annals of Mathematics and Artificial Intelligence*, 77(3–4), 303–316. Springer.
- Biró, P., Manlove, D. F., & McBride, I. (2014). The hospitals/residents problem with couples: Complexity and integer programming models. In *Proceedings of the International Symposium on Experimental Algorithms* (p. 10–21). Springer.
- Biró, P., Manlove, D. F., & Mittal, S. (2010). Size versus stability in the marriage problem. *Theoretical Computer Science*, 411(16–18), 1828–1841.
- Bronfman, S., Hassidim, A., Afek, A., Romm, A., Shreberk, R., Hassidim, A., & Massler, A. (2015). Assigning Israeli medical graduates to internships. *Israel Journal of Health Policy Research*, 4(6).
- Budish, E., Cachon, G. P., Kessler, J. B., & Othman, A. (2016). Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 65(2), 314–336.
- Cechlárová, K., Eirinakis, P., Fleiner, T., Magos, D., Manlove, D. F., Mourtos, I., Oceláková, E., & Rastegari, B. (2016). Pareto optimal matchings in many-to-many markets with ties. *Theory of Computing Systems*, 59(4), 700–721.
- Cechlárová, K., & Fleiner, T. (2005). On a generalization of the stable roommates problem. *ACM Transactions on Algorithms (TALG)*, 1(1), 143–156. ACM
- Chen, Y., & Sönmez, T. (2006). School choice: An experimental study. *Journal of Economic Theory*, 127(1), 202–231.
- Chiarandini, M., Fagerberg, R., & Gualandi, S. (2017). Handling preferences in student-project allocation. *Annals of Operations Research*, 1–40. doi:10.1007/s10479-017-2710-1.
- Combe, J., Tercieux, O., & Terrier, C. (2018). *The design of teacher assignment: Theory and evidence*. Mimeo.
- de Haan, M., Gautier, P. A., Oosterbeek, H., & Van der Klaauw, B. (2018). The performance of school assignment mechanisms in practice. Working Paper, Delacrétaz, D., Kominers, S. D., & Teytelboym, A. (2016). *Refugee resettlement*. Mimeo.
- Diebold, F., & Bichler, M. (2017). Matching with indifferences: A comparison of algorithms in the context of course allocation. *European Journal of Operational Research*, 260(1), 268–282.
- Erdil, A., & Ergin, H. (2008). What's the matter with tie-breaking? Improving efficiency in school choice. *American Economic Review*, 98(3), 669–689.
- Firat, M., Briskorn, D., & Laugier, A. (2016). A branch-and-price algorithm for stable workforce assignments with hierarchical skills. *European Journal of Operational Research*, 251(2), 676–685.
- Fleiner, T. (2003). A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research*, 28(1), 103–126. INFORMS.

- Gale, D. (1960). *The theory of linear economic models*. New York: McGraw-Hill, Inc.
- Gale, D., & Shapley, L. (1962). College admissions and the stability of marriage. *American Mathematical Monthly*, 69, 9–15.
- Gale, D., & Sotomayor, M. (1985). Some remarks on the stable matching problem. *Discrete Applied Mathematics*, 11(3), 223–232.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. San Francisco, CA: Freeman.
- Garg, N., Kavitha, T., Kumar, A., Mehlhorn, K., & Mestre, J. (2010). Assigning papers to referees. *Algorithmica*, 58(1), 119–136.
- Haeringer, G. (2017). *Market design: Auctions and matching*. Cambridge, MA: MIT Press.
- Harless, P. (2019). *Immediate acceptance with or without skips? Comparing school assignment procedures*.
- Hatfield, J. W., & Milgrom, P. R. (2005). Matching with contracts. *American Economic Review*, 95(4), 913–935. American Economic Association
- Irving, R. W., Leather, P., & Gusfield, D. (1987). An efficient algorithm for the “optimal” stable marriage. *Journal of the ACM (JACM)*, 34(3), 532–543.
- Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2, 83–97.
- Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11), 3257–3287.
- Manlove, D. F. (2013). *Algorithmics of matching under preferences*. World Scientific Publishing Company.
- Manlove, D. F., Irving, R. W., Iwama, K., Miyazaki, S., & Morita, Y. (2002). Hard variants of stable marriage. *Theoretical Computer Science*, 276, 261–279.
- Moraga, J. F. H., & Rapoport, H. (2014). Tradable immigration quotas. *Journal of Public Economics*, 115, 94–108.
- Mossaad, N., Ferwerda, J., Lawrence, D., Weinstein, J. M., & Hainmueller, J. (2018). Determinants of refugee naturalization in the United States. *Proceedings of the National Academy of Sciences*, 115(37), 9175–9180.
- Nguyen, T., Nguyen, H., & Teytelboym, A. (2019). Stability in matching markets with complex constraints. In *Proceedings of the ACM Conference on Economics and Computation* (p. 61). ACM.
- Othman, A., Sandholm, T., & Budish, E. (2010). Finding approximate competitive equilibria: Efficient and fair course allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems. 1, 873–880
- Perach, N., Polak, J., & Rothblum, U. G. (2008). A stable matching model with an entrance criterion applied to the assignment of students to dormitories at the technion. *International Journal of Game Theory*, 36, 519–535.
- Roth, A. E. (1984). The evolution of the labor market for medical interns and residents: A case study in Game Theory. *Journal of Political Economy*, 92(6), 991–1016.
- Roth, A. E. (1986). On the allocation of residents to rural hospitals: A general property of two-sided matching markets. *Econometrica*, 425–427.
- Roth, A. E., & Peranson, E. (1999). The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89(4), 748–780.
- Roth, A. E., & Postlewaite, A. (1977). Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, 4(2), 131–137.
- Roth, A. E., Sönmez, T., & Ünver, M. U. (2004). Kidney exchange. *The Quarterly Journal of Economics*, 119(2), 457–488. MIT Press.
- Saban, D., & Sethuraman, J. (2015). The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research*, 40(4), 1005–1014.
- Scarf, H. E. (1967). The core of an n person game. *Econometrica*, 69, 35–50.
- Shapley, L., & Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, 1, 23–27.
- Shi, P. (2015). Guiding school-choice reform through novel applications of operations research. *Interfaces*, 45(2), 117–132.
- Shi, P. (2016). *Assortment planning in school choice*. Mimeo.
- Shorrer, R. I., & Sívágó, S. (2018). Obvious mistakes in a strategically simple college-admissions environment. Working paper.
- Sönmez, T., & Yenmez, M. B. (2019). *Affirmative action in India via vertical and horizontal reservations*. Mimeo.
- Trapp, A. C., Teytelboym, A., Martinello, A., Andersson, T., & Ahani, N. (2018). Placement optimization in refugee resettlement. Working Paper 2018:23. Lund University.
- Veski, A., Biró, P., Pöder, K., & Lauri, T. (2017). Efficiency and fair access in kindergarten allocation and policy design. *Journal of Mechanism and Institution Design*, 2(1), 57–104.
- Yang, W., Giampapa, J. A., & Sycara, K. (2003). Two-sided matching for the US Navy detailing process with market complication. Technical Report CMU-RI-TR-03-49, Robotics Institute, Carnegie-Mellon University.