On Matching Problems in Large Settings

by

Ishan Agarwal

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Computer Science

New York University

September, 2023

Professor Richard Cole

DEDICATION

For my mother and father, who have always inspired me to never give up, even during the most difficult times. Also to my advisor and mentor, Richard Cole; without his help and continuous guidance, this would have never been possible.

Finally to the wonderful people at NYU, especially Oded Regev, from whom I have learnt so much, and to my closest friends, who have been a bulwark of sanity in a world filled with madness.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Richard Cole, for his role in both inspiring as well as tackling the questions considered in this work. His guidance has been invaluable in developing my ability not only to understand and approach challenging research questions but also to present my ideas cogently.

I would also like to thank the committee members, especially Oded Regev and Subhash Khot, for their feedback throughout this project.

Finally, I would like to thank my parents for their unwavering belief in me, and for their moral support throughout each day of this journey.

ABSTRACT

Matching problems arise in several settings in practice and have been a longstanding subject of theoretical analysis. Typically, the settings of interest involve a large number of agents. We further the study of matching problems in two settings: the stable matching setting, which has been studied since the seminal work of Gale and Shapley, and a setting where agents' values to prospective partners degrade over time, leading them to have to balance the trade-off between searching for a better partner versus deciding to match.

In the stable matching setting, we extend a line of research that seeks to explain the dichotomy between the fact that Gale and Shapley's Deferred Acceptance algorithm seems to work well in practice, even when agents only submit a short list of prospective partners to the centralized matching algorithm, and the fact that if the agents' preferences are allowed to be arbitrary, complete lists of all agents' preferences are needed in order to guarantee a stable matching. To this end, we consider probabilistically generated preference lists and we show that under fairly general assumptions and in a variety of models, with high probability, short lists of prospective partners, namely length $\Theta(\log n)$ instead of n, suffice for most of the agents. We prove our bounds are tight up to constant factors. Furthermore, we construct a simple set of $\Theta(\log n)$ possible matches per agent for almost all agents and demonstrate (in the form of an approximate equilibrium result) that they can afford to restrict their proposals to this set, while incurring only a small loss in utility.

In the time discounted utilities setting, we consider a dynamic matching market, and study how agents should balance accepting a proposed match with the cost of continuing their search. Our model has two new features: finite agent lifetimes with linear loss in utility over time, and a discrete population model, aspects which are underexplored in the literature. We quantify how well the agents can do by providing upper and lower bounds on the collective losses of the agents, with a polynomially small failure probability, where the notion of loss is with respect to a plausible baseline we define. These bounds are also tight up to constant factors.

In both settings, we complement our theoretical results with numerical simulations.

Contents

List of Figures

List of Tables

List of Appendices

1 Introduction

1.1 A Perspective on Algorithmic Game Theory

Algorithmic game theory is a field that combines concepts from computer science and economics to study strategic interactions among agents. People interacting among themselves, either with or without the mediation of a centralised third party, are typically modeled as rational agents, and the rules governing their interaction are modeled by a mechanism. In the field of mechanism design, a mechanism refers to an algorithm that takes inputs from the participants in the game or economic system and outputs an outcome.

At a very high level, in this setting, there are two very different directions of study:

- **Designing good mechanisms**

  Typically there is a tension between the desires of each agent and the sort of outcome the mechanism designer wants. To resolve this, a common goal is to construct mechanisms that align the individual incentives of the participants with the desired collective outcome. The alignment is often achieved by creating mechanisms in which the agents have no incentive to misreport their preferences, thus removing the need for participants to strategize about what to report. Standard desirable outcomes include satisfying certain fairness properties, or being optimal with respect to some overall measure of quality (see [Finocchiaro et al. 2021] for a detailed discussion of this literature). Finally, the mechanism should ideally not be too computationally complex (the importance of this requirement has been recognized from very early on: see [Suppes 1958]).

Thus mechanism design focuses on the question of what makes a mechanism ‘good’ and seeks to design ’good’ mechanisms in a variety of settings. In the opposite direction there is a corresponding literature of work showing hardness bounds and impossibility results that set out the limits on what ‘good’ properties can possibly be achieved by any mechanism. See [Papadimitriou et al. 2008; Dobzinski et al. 2022; Feigenbaum et al. 2003] for examples.

- **Explaining agent behaviour and why existing mechanisms work well**

  Several mechanisms that have poor worst case guarantees seem to work quite well in practice. We will see an example of this in Chapter 2, with regards to the stable matching problem. The explanation may be that in reality the agents participating in the mechanism do not have arbitrary preferences but rather there are specific features of real world settings that enable these mechanisms to work well. We can seek to characterize such features.

  Furthermore people in the real world do not engage in arbitrary strategies when interacting with others. We can seek to explain people’s behaviour and/or provide useful advice in the form of simple strategies that agents can follow that will yield close to optimal outcomes.

For an excellent introduction to the literature exploring both of the important directions of reasearch in algorithmic game theory generally, along with examples, we refer the reader to [Roughgarden 2008]. The questions this thesis studies are all in the domain of the second topic.

Using ideas from computer science to explain the behaviour of interacting agents, is a long-standing line of work. [Friedman and Shenker 1998] points out the importance of rethinking notions of rationality and equilibrium concepts in an effort to come up with more compelling models of agent behaviour. See ([Blum et al. 2008; Babaioff et al. 2006; Fabrikant and Papadimitriou 2008; Nisan et al. 2008]) for examples of work that couple worst case guarantees with novel behavioural models. Even more broadly, there is a burgeoning literature of work seeking to explain results from economics and the social sciences by constructing appropriate mathematical models and applying ideas from theoretical computer science. One early shining example of this is Jon Kleinberg's work aimed at explaining the famous 'small world phenomenon' from the social sciences literature ([Kleinberg 2000b]).

The small-world phenomenon refers to the idea that we are all connected through short chains of acquaintances. It is a fundamental concept in social networks, which can be understood as the presence of numerous short paths in a graph representing people, where edges represent acquaintances between individuals. This observation was famously demonstrated by social psychologist Stanley Milgram in the 1960s by means of experiments. Milgram ([Milgram 1967]) conducted a study where participants in Kansas and Nebraska were asked to forward a letter to a "target person" near Boston, with the condition that each person could forward the letter to just one acquaintance. The results showed that the median length of completed chains was six, since popularized via the term "six degrees of separation."

Kleinberg raised two important questions:

- Why should there exist short chains of acquaintances linking together arbitrary pairs of strangers?
- Why should arbitrary pairs of strangers be able to find short chains of acquaintances that link them together?

The first question had already been partly addressed. One of Kleinberg's insights was the important distinction between the two questions. He proposed a relatively simple family of random network models and showed that for exactly one of these models, there was an efficient decentralized algorithm to find short paths ([Kleinberg 2000a, 2001, 2006]).

As Kleinberg observed:

> The ability to construct a searchable network in this way has since proved useful in the design of peer-to-peer file-sharing systems on the Internet, where content must be found by nodes consulting one another in a decentralized fashion. In other words, nodes executing these look-up protocols are behaving very much like participants in the Milgram experiments – a striking illustration of the way in which the computational and social sciences can inform one another, and the way in which mathematical models in the computational world turn into design principles with remarkable ease.
>
> - Jon Kleinberg, *The small-world phenomenon and decentralized search* [*Kleinberg 2004*]

The questions considered in this thesis have a similar flavour, insofar as we too seek to understand and explain agent behaviour and characterize what allows mechanisms to work surprisingly well in real world settings. Our focus is on matching problems, a topic that comes up in numerous real world scenarios: matching students to colleges, employees to employers and dating applications are just a few salient examples.

In Chapter 2 we consider the Deferred Acceptance (DA) algorithm of Gale and Shapley [Gale and Shapley 1962]. When agents do not list all possible partners in order of preference, there is, in general, no guarantee that the agents will all be matched. Moreover the DA algorithm is not incentive compatible for all agents: agents may be able to get a better partner by misreporting their preferences. In reality, agents often do not list their preferences among all possible partners: eliciting this many preferences is generally too time consuming or expensive. In fact, people's preferences are correlated, not arbitrary. We model this correlation and show that in this setting the DA mechanism, with high probability, matches all agents and that the agents obtain reasonably good matches. Furthermore, we are able to show that agents do not have much to gain by misreporting their preferences. This provides a simple intuitive answer to the question: where should (as an example) medical residents apply to for residency positions. The answer turns out to be: apply to a range of hospitals ranked about as highly as you are publicly perceived to be ranked among the pool of applicants, and of the hospitals within this range, apply to your favourite ones.

In Chapter 3, we explore how agents should behave when faced with the task of finding a match when they have to balance the trade-off between exploring for better matches and accepting a match so as to enjoy the match for a longer period of time. We consider a model with fixed finite agent lifetimes, and the utility derived by the agents is the product of the quality of the match and the minimum of the two matched agents' lifetimes. While this is quite a crude approximation of any real world setting, it still captures important qualitative features of agent behaviour: for example, as might be expected, the agents should get less picky with time. Even in our setting, finding an exactly optimal strategy for each agent is quite hard and our investigation suggests these strategies are quite complex. However, we are able to analyze a much simpler strategy which still yields reasonable outcomes.

1.2 Types of Matching Problems

Matching problems, broadly speaking, involve scenarios where items or agents need to be paired up in some way. Such pairing up tasks are common in practice and matching problems have been widely studied in the mathematics, computer science and economics literature.

One broad type of problem consists of purely combinatorial matching problems: typical examples being the task of finding or counting maximal or maximum matchings in graphs (for a detailed survey see [Biggs 1988]; [Avis 1983] surveys matchings involving points in Euclidean space). These kinds of problems and their variants remain an important topic of study. However, we will focus on settings where the entities being matched have some notion of preferences between various possible match options.

Matching problems of this kind can themselves be categorized across several dimensions.

1. The matching problem may be one-sided or two-sided. In a one-sided matching, agents are matched to items; each agent has an ordering of the items based on their own preferences, but the items have no associated preferences. This setting is often called the assignment or allocation problem and is used to model several real world settings. A typical example would be a market where items must be allocated to bidders for a price [Abdulkadiroglu and Sönmez 2013]. There are also allocation problems without money, such as the task of assigning classes to students, where classes have capacities and there is the further constraint that the class allocations for each student must be non-overlapping [Diebold et al. 2014; Diebold and Bichler 2017].

   Two-sided matchings refer to settings where both sides being matched have preference orderings over the other side. Common real world examples include dating applications [Abadi and Prabhakar 2017], matching employers and employees as well as ride-share mechanisms [Chau et al. 2020]. Once again, each of these applications may come with its own set of additional constraints. Our work focuses on two-sided matching problems.

2. The preferences of the agents may be ordinal or cardinal. In the ordinal preference setting, the agents each have an ordered list of their potential matches, while in the case of cardinal preferences the agents have an underlying utility value associated with their possible partners. From a modelling perspective, these two types of preferences both have pros and cons. While ordinal preferences are easier to elicit, cardinal utilities are more expressive, enabling an agent to report not only if they prefer one good/partner to another but also by how much. This can allow for better measures of the quality of the matching. In both Chapters 2 and 3 we work with cardinal utility models.

3. The matching algorithm may seek to optimize a variety of different measures of the quality of the matching. In one-sided matching, the mechanism assigning items to bidders may be trying to maximize total utility or it may instead seek to optimize some metric of fairness

(especially commonly studied in allocation problems without money [Amanatidis et al. 2017]). In two-sided matching problems, one solution concept that has been explored in detail in the literature is the notion of stable matching. We will discuss this in detail in Chapter 2.

4. The matching may be performed by some centralized mechanism or the setting may be that of a dynamic matching market where the agents themselves encounter one another and decide on whom to match with. A real world example of the former would be the US medical residency program, where doctors and hospitals submit their preferences to a centralized matching mechanism that runs a version of the deferred acceptance algorithm and outputs a stable matching. Our work in Chapter 2 examines this setting. A less centralized process is considered in Chapter 3, which examines a dynamic matching market where agents are paired up at random and they can either choose to match or not match. Often in real-world applications there is a mix of both centralized and individual decision making: for example a dating application might suggest potential matches, not wholly at random but rather based on expressed preferences, and then the agents on both sides make the choice to match or not.

5. The setting may be online or offline. The matching mechanism may receive the information about the agents/items being matched all at once (the offline setting), or it may receive such information in an online manner. A well-studied example studied in the online setting is the item allocation problem: the mechanism sees one item arrive at a time and it only sees the agents' bids for an item once the item arrives (see [Mehta et al. 2013] for a detailed survey of one-sided online matching problems). The mechanism then has to make a (typically irrevocable) assignment of the item to some bidder before it gets to see the next item. Another well-studied problem in the online setting is the online bipartite matching problem, where algorithms have to decide whether to add edges to the matching being

constructed when each edge arrives. In models with online arrivals, the items may arrive in a random or adversarial order.

The work presented in Chapter 2 considers the stable matching problem, while Chapter 3 considers a dynamic matching market with agent utilities diminishing over time.

2 Stable Matching: Choosing Which Proposals to Make

2.1 Introduction

Consider a doctor applying for residency positions. Where should she apply? To the very top programs for her specialty? Or to those where she believes she has a reasonable chance of success (if these differ)? And if the latter, how does she identify them? We study these questions in the context of Gale and Shapley's Deferred Acceptance (DA) algorithm [Gale and Shapley 1962].

We will begin with a brief discussion of the Stable Matching problem and the classical Deferred Acceptance algorithm, since these will be of central importance throughout this chapter.

2.1.1 Stable Matching and the Deferred Acceptance (DA) Algorithm

The stable matching problem is classically stated in terms of matching men and women. Let M be a set of n men and W a set of n women. Each man m has an ordered list of women that represents his preferences, i.e. if a woman w comes before a woman w' in m's list, then m would prefer matching with w rather than w'. The position of a woman w in this list is called m's ranking of w. Similarly each woman w has a ranking of the men[1]. The stable matching task is to pair

[1] Throughout this chapter, we assume that each man m (woman w) ranks all the possible women (men), i.e. m's (w's) preference list is complete.

ALGORITHM 1: Woman Proposing Deferred Acceptance (DA) Algorithm.

Initially, all the men and women are unmatched.
while *some woman w with a non-empty preference list is unmatched* **do**
  let m be the first man on her preference list;
  if *m is currently unmatched* **then**
    tentatively match w to m.
  end
  if *m is currently matched to w', and m prefers w to w'* **then**
    make w' unmatched and tentatively match w to m.
  end
  remove m from w's preference list
end

(match) the men and women in such a way that no two people prefer each other to their assigned partners. More formally:

Definition 2.1 (Matching)**.** A matching is a pairing of the agents in M with the agents in W. It comprises a bijective function μ from M to W, and its inverse $\nu = \mu^{-1}$, which is a bijective function from W to M.

Definition 2.2 (Blocking pair)**.** A matching μ has a blocking pair (m, w) if and only if:

1. m and w are not matched: $\mu(m) \neq w$.

2. m prefers w to his current match $\mu(m)$.

3. w prefers m to her current match $\nu(w)$.

Definition 2.3 (Stable matching)**.** A matching μ is stable if it has no blocking pair.

Gale and Shapley [Gale and Shapley 1962] proposed the seminal deferred acceptance (DA) algorithm for the stable matching problem. We present the woman-proposing DA algorithm (Algorithm 1); the man-proposing DA is symmetric.

Gale and Shapley's work was motivated by issues they observed in the US college admission process. [Gale et al. 2001] contains a discussion on the process of the development of the algorithm. Gale states that the starting point of their investigation was a New York Times article on

10th September 1960 in which a reporter observing the undergraduate admissions process at Yale noted that:

> “the admissions men very often have no way of discovering how many other colleges each applicant is trying for, nor have they any way of knowing how many students they decide to admit actually intend to come to their college"

With the consequence that:

> “the admissions officer may end up "discovering that he has acceptances from a freshman class either half as large or twice as large as the school has room for. ...because of all the guess work, one would expect the final allocation of applicants to colleges would be highly “non-optimal", so the first problem was to pin down precisely the nature of these “non-optimalities". With this in mind, I decided to look first at the special case where each college has a quota of one."
>
> \- David Gale, *The two-sided matching problem. Origin, development and current issues [Gale et al. 2001]*

The idea of stability of the form we described above (absence of blocking pairs) is quite intuitive, as evidenced by the fact that this notion of stability and indeed deferred acceptance type algorithms have been informally discovered and implemented repeated times independently. The celebrated result of Gale and Shapley formalized these intuitions and proved the correctness of the algorithm: showing that with complete preferences it is guaranteed to find a stable matching.

The need for a centralized matching mechanism had become evident in various other settings as well. One striking example was the case of medical internships. Internships have been a form of postgraduate medical education since around 1900. For hospitals this provided a source of relatively cheap labour; thus there was significant competition among hospitals for interns. As a consequence hospitals tried to set the date for the binding agreements earlier than their competitors. By 1944 dates of appointments were up to two full years before the internship was actually to begin. Students waited for offers from preferable positions, and hospitals got last minutes rejections. Ultimately a centralized clearinghouse was proposed: instead of hospitals making individual offers and students responding, students and programs would submit rank order list to indicate their preferences. In 1950–1951 a trial run of the centralized algorithm was held and in 1951–1952 the new NIMP algorithm was implemented which was very similar to what we now know as the DA algorithm. This program gets very high levels of voluntary participation up to the present day. For a detailed discussion of this history, see [Roth 1984].

The matching system based on DA has even faced legal challenges. In 2002 law firms brought an antitrust suit against the matching system on the grounds that it was a conspiracy to hold down wages for medical residents. The counter argument was that the algorithms based on DA have been shown to be robust and useful in a wide variety of contexts, including the medical residency match. This latter perspective prevailed, and the use of the Deferred Acceptance algorithm has been explicitly recognized as part of pro-competitive market mechanism in American law. For a detailed survey of the literature on stable-matching problem and the DA algorithm, including the history of use of the DA algorithm in the US medical residency match context, see [Roth 2008].

Nowadays, DA algorithm is widely used to compute matchings in a wide variety of real-world applications: the National Residency Matching Program (NRMP), which matches future residents to hospital programs [Roth and Peranson 1999]; university admissions programs which match students to programs, e.g. in Chile [Rios et al. 2021], school choice programs, e.g. for placement in New York City's high schools [Abdulkadiroğlu et al. 2005], the Israeli psychology Masters match [Hassidim et al. 2017], and many others (e.g. [Gonczarowski et al. 2019]).

Gale and Shapley's work sparked a great deal of work on the stable matching problem and more generally on various approaches to the two sided matching problem with preferences. See [Iwama and Miyazaki 2008] for a detailed review of various extensions and generalisations of the stable matching problem and [Ren et al. 2021] for a survey on matching algorithms. [Biró and Klijn 2013] also reviews the literature on the stable matching problem, including the generalisation to the many-to-one matching setting, before considering the problem with additional constraints. For a detailed discussion of applications, see [Biró 2017]. [Hakimov and Kübler 2019] provides a discussion of experimental results. In addition there are numerous books on the topic: see [Gusfield and Irving 1989] for an overview, [Roth and Sotomayor 1992; Manlove 2013] for an algorithmic perspective and [Knuth 1996] for a discussion of how the stable matching problem can be related to other combinatorial problems.

The following facts about the DA algorithm are well known and we will use them freely throughout this chapter.

Lemma 2.4. *DA terminates and outputs a stable matching.*

Proof. The DA algorithm always terminates since no woman proposes twice to the same man. During an execution of the woman-proposing DA algorithm, a man, once matched, remains matched. Suppose the algorithm ends with an unmatched woman w, and hence an unmatched man m. But w would have proposed to m at some point, so m would not be unmatched. Thus when the algorithm terminates, every man and woman is matched.

Suppose the matching resulting from DA is not stable. In the output matching suppose w_1 is matched with m_1 and w_2 is matched with m_2. Suppose (w_1, m_2) form a blocking pair. Then w_1 prefers m_2 to m_1. This means that, according to the algorithm, she must have already proposed to m_2 and been rejected by him. This means that m_2 strictly prefers w_2 to w_1, since a man only improves his match over time. Hence w_1 and m_2 cannot form a blocking pair. □

The following result is an immediate consequence of Lemma 2.4.

Corollary 2.5. *For any given set of preferences of the men and women, a stable matching always exists.*

Lemma 2.6. *Woman-proposing DA is woman-optimal, i.e. each woman is matched with the best partner she could be matched with in any stable matching.*

Proof. The woman-proposing DA algorithm is under-specified: at each step we have a choice regarding which unmatched woman is chosen to propose. Consider any fixed execution of the woman-proposing DA algorithm. Let $R = \{(w, m)\}$ denote the set of pairs (m, w) such that m rejects w at some point in this fixed execution of the woman-proposing algorithm. Since each woman systematically works her way down her preference list, if w is matched to m at the conclusion of the algorithm, then $(w, m') \in R$ for every m' that w prefers to m'. Thus, the following claim would imply the theorem: for every $(w, m) \in R$, no stable matching pairs up w and m.

We prove the claim by induction on the number of iterations of the algorithm. Initially, $R = \phi$ and the claim is true. For the inductive step, consider an iteration of the algorithm in which m rejects w in favor of w'. Thus one of w, w' proposed to m in this iteration. Since w' makes proposals in the order specified by her preference list, for every m' that w' prefers to m, (w', m') is already in the current set R of rejected proposals. By the inductive hypothesis, in every stable matching, w' is paired with m or someone she prefers less. Since m prefers w' to w, and w' prefers m to anyone else she might be matched to in a stable matching, there is no stable matching that pairs w with m since otherwise (w', m) would form a blocking pair. □

The following result is an immediate consequence of Lemma 2.6.

Corollary 2.7. *The stable matching generated by DA is independent of the order in which the unmatched agents on the proposing side are processed.*

Lemma 2.8. *Woman-proposing DA is man-pessimal, i.e. each man is matched with the worst partner he could be matched with in any stable matching.*

Proof. Let P be the woman-optimal matching, and suppose that P was not man-pessimal. By definition this would mean that there is some man m who could be worse off with some other stable matching P'. Let's say that P pairs m with w, and P' pairs m with w'. By assumption m likes w better than w'. However, since P' pairs m to w', it must pair w with some other man, m'. However, since P is assumed to be woman-optimal, m is a man w prefers to any other man she can be matched with in any stable matching. In particular, w likes m better than m'. But this means that m and w form a blocking pair for P', which contradicts the stability of P'. Thus a woman-optimal matching must be man-pessimal and thus the result now follows by Lemma 2.6 □

The DA algorithm is incentive compatible for the proposing side. That is, no individual agent on the proposing side can gain a more preferred partner by misreporting their preference list.

Theorem 2.9. *[Dubins and Freedman 1981; Roth 1982] In the game induced by the man-proposing deferred acceptance algorithm, in which each player states a preference list, it is a weakly dominant strategy for each man to state his true preferences.*

However the same incentive compatibility properties do not apply to the side receiving proposals. In fact there is no mechanism to output stable matchings that is incentive compatible for every participating agent.

Theorem 2.10. *[Roth and Sotomayor 1990] When any stable mechanism is applied to a marriage market in which preferences are strict and there is more than one stable matching, then at least one agent can profitably misrepresent his or her preferences, assuming the others tell the truth. (This agent can misrepresent so that in every stable matching under the mis-stated preferences, they are matched to their most preferred achievable partner under the true preferences.)*

2.1.2 Stable matching with short preference lists

It is well-known that in DA the optimal strategy for the proposing side is to list their choices in order of preference. However, this does not address *which* choices to list. This issue has been relatively overlooked in the literature.

Recall that each agent provides the mechanism a list of its possible matches in preference order, including the possibility of "no match" as one of its preferences. These mechanisms promise that the output will be a stable matching with respect to the submitted preference lists. In practice, submitted preference lists are relatively short. This may be directly imposed by the mechanism or could be a reflection of the costs—for example, in time or money—of determining these preferences. Note that a short preference list is implicitly stating that the next preference after the listed ones is "no match".

Thus it is important to understand the impact of short preference lists. Roth and Peranson observed that the NRMP data showed that preference lists were short compared to the number of programs and that these preferences yielded a single stable partner for most participants; we note that this single stable partner could be the "no match" choice, and in fact this is the outcome for a constant fraction of the participants. They also confirmed this theoretically for the simplest model of uncorrelated random preferences; namely that with the preference lists truncated to the top O(1) preferences, almost all agents have a unique stable partner. Subsequently, in [Immorlica and Mahdian 2015] the same result was obtained in the more general popularity model which allows for correlations among different agents' preferences; in their model, the first side—men—can have arbitrary preferences; on the second side—women—preferences are selected by weighted random choices, the weights representing the "popularity" of the different choices. These results were further extended by Kojima and Parthak in [Kojima and Pathak 2009].

The popularity model does not capture behavior in settings where bounds on the number of proposals lead to proposals being made to plausible partners, i.e. partners with whom one has a realistic chance of matching. One way to capture such settings is by way of tiers [Ashlagi et al. 2019], also known as block correlation [Coles et al. 2013]. Here agents on each side are partitioned into tiers, with all agents in a higher tier preferred to agents in a lower tier, and with uniformly random preferences within a tier. Tiers on the two sides may have different sizes. If we assign tiers successive intervals of ranks equal to their size, then, in any stable matching, the only matches will be between agents in tiers whose rank intervals overlap.

A more nuanced way of achieving these types of preferences bases agent preferences on cardinal utilities; for each side, these utilities are functions of an underlying common assessment of the other side, together with idiosyncratic individual adjustments for the agents on the other side. These include the separable utilities defined by Ashlagi, Braverman, Kanoria and Shi in [Ashlagi et al. 2019], and another class of utilities introduced by Lee in [Lee 2016]. This last model will be the focus of our study.

To make this more concrete, we review a simple special case of Lee's model, the *linear separable model*. Suppose that there are n men and n women seeking to match with each other. Each man m has a public rating r_m, a uniform random draw from $[0, 1]$. These ratings can be viewed as the women's joint common assessment of the men. In addition, each woman w has an individual adjustment, which we call a score, $s_w(m)$ for man m, again a uniform random draw from $[0, 1]$. All the draws are independent. Woman w's utility for man m is given by $\frac{1}{2}[r_m + s_w(m)]$; her full preference list has the men in decreasing utility order. The men's utilities are defined similarly.

Lee stated that rather than being assumed, short preference lists should arise from the model; this appears to have been a motivation for the model he introduced. A natural first step would be to show that for some or all stable matchings, the utility of each agent can be well-predicted, for this would then allow the agents to limit themselves to the proposals achieving such a utility. Lee proved an approximate version of this statement, namely that with high probability (w.h.p., for short) most agents obtain utility within a small ε of an easily-computed individual benchmark. Unfortunately, it is not clear this is enough to allow agents to limit their proposals as just indicated.

2.1.3 Our Contribution

Our work seeks to resolve this issue. We obtain the following results. Note that in these results, when we refer to the bottommost agents, we mean when ordered by decreasing public rating. Also, we let the term loss mean the difference between an agent's benchmark utility and their achieved utility.

1. We show that in the linearly separable model, for any constant $c > 0$, with probability $1 - 1/n^c$, in every stable matching, apart from a sub-constant σ fraction of the bottommost agents, *all* the other agents obtain utility equal to an easily-computed individual benchmark $\pm\varepsilon$, where ε is also sub-constant.

   We show that both $\sigma, \varepsilon = \widetilde{\Theta}(n^{-1/3})$.[2] As we will see, this implies, w.h.p., that for all the agents other than the bottommost σ fraction, each agent has $\Theta(\ln n)$ possible edges (proposals) that could be in any stable matching, namely the proposals that provide both agents utility within ε of their benchmark. Furthermore, we show our bound is tight: with fairly high probability, there is no matching, let alone stable matching, providing every agent a partner if the values of ε and σ are reduced by a suitable constant factor.

   An interesting consequence of this lower bound on the agents' utilities is that the agents can readily identify a moderate sized subset of the edge set to which they can safely restrict their applications. More precisely, any woman w outside the bottommost σ fraction, knowing only her own public rating, the public ratings of the men, and her own private score for each man, can determine a preference list of length $\widetilde{\Theta}(n^{1/3})$ which, w.h.p, will yield the same result as her true full-length list. Our analysis also shows that if w obtained the men's private scores for these proposals, then w.h.p. she could safely limit herself to a length $O(\ln n)$ preference list.

2. The above bounds apply not only to the linearly separable model, but to a significantly more general bounded derivative model (in which derivatives of the utility functions are bounded).

3. The result also immediately extends to settings with unequal numbers of men and women. Essentially, our analysis shows that the loss for an agent is small if there is a σ fraction of agents of lower rank on the opposite side. Thus even on the longer side, w.h.p., the topmost

[2]The $\widetilde{\Theta}(\cdot)$ notation means up to a poly-logarithmic term; here $\sigma, \varepsilon = \Theta((n/\ln n)^{-1/3})$.

$n(1-\sigma)$ agents all obtain utility close to their benchmark, where n is the size of the shorter side. This limits the "stark effect of competition" [Ashlagi et al. 2017]—namely that the agents on the longer side are significantly worse off—to a lower portion of the agents on the longer side.

4. The result extends to the many-to-one setting, in which agents on one side seek multiple matches. Our results are given w.r.t. a parameter d, the number of matches that each agent on the "many" side desires. For simplicity, we assume this parameter is the same for all these agents. In fact, we analyze a more general many-to-many setting.

5. A weaker result with arbitrarily small $\sigma, \varepsilon = \Theta(1)$ holds when there is no restriction on the derivatives of the utility functions, which we call the general values model. Again, we show this bound cannot be improved in general. This setting is essentially the general setting considered by Lee [Lee 2016]. He had shown there was a σ fraction of agents who might suffer larger losses; our bound identifies this σ fraction of agents as the bottommost agents.

6. In the bounded derivative model, with slightly stronger constraints on the derivatives, we also show the existence of an ϵ-Bayes-Nash equilibrium in which no agent proposes more than $O(\ln^2 n)$ times and all but the bottommost $O((\ln n/n)^{1/3})$ fraction of the agents make only the $O(\ln n)$ proposals identified in (1) above. Here $\varepsilon = \Theta(\ln n/n^{1/3})$.

These results all follow from a lemma showing that, w.h.p., each non-bottommost agent has at most a small loss. In turn, the proof of this lemma relies on a new technique which sidesteps the conditioning inherent to runs of DA in these settings.

Experimental results Much prior work has been concerned with preference lists that have a constant bound on their length. For moderate values of n, say $n \in [10^3, 10^6]$, $\ln n$ is quite small, so our $\Theta(\ln n)$ bound may or may not be sufficiently small in practice for this range of n. What matters are the actual constants hidden by the Θ notation, which our analysis does not fully determine. To help resolve this, we conducted a variety of simulation experiments.

We have also considered how to select the agents to include in the preference lists, when seeking to maintain a constant bound on their lengths, namely a bound that, for the values of n we considered, was smaller than the $\Theta(\ln n)$ bound determined by the above simulations; again, our investigation was experimental.

2.1.4 Other Related work

The random preference model was introduced by Knuth [Knuth 1976] (for a version in English see [Knuth 1996]), and subsequently extensively analyzed [Pittel 1989; Knuth et al. 1990; Pittel 1992; Mertens 2005; Pittel et al. 2008; Pittel 2019; Kupfer 2020]. In this model, each agent's preferences are an independent uniform random permutation of the agents on the other side. An important observation was that when running the DA algorithm, the proposing side obtained a match of rank $\Theta(\ln n)$ on the average, while on the other side the matches had rank $\Theta(n/\ln n)$.

A recent and unexpected observation in [Ashlagi et al. 2017] was the "stark effect of competition": that in the random preferences model the short side, whether it was the proposing side or not, was the one to enjoy the $\Theta(\ln n)$ rank matches. Subsequent work showed that this effect disappeared with short preference lists in a natural modification of the random preferences model [Kanoria et al. 2021]. Our work suggests yet another explanation for why this effect may not be present: it does not require that short preference lists be imposed as an external constraint, but rather that the preference model generates few edges that might ever be in a stable matching.

The number of edges present in any stable matching has also been examined for a variety of settings. When preference lists are uniform the expected number of stable pairs is $\Theta(n \ln n)$ [Pittel 1992]; when they are arbitrary on one side and uniform on the other side, the expected number is $O(n \ln n)$ [Knuth et al. 1990]. This result continues to hold when preference lists are arbitrary on the men's side and are generated from general popularities on the women's side [Gimbert et al.

2019]. Our analysis shows that in the linear separable model (and more generally in the bounded derivative setting) the expected number of stable pairs is also $O(n \ln n)$.

Another important issue is the amount of communication needed to identify who to place on one's preference lists when they have bounded length. In general, the cost is $\Omega(n)$ per agent (in an n agent market) [Gonczarowski et al. 2015], but in the already-mentioned separable model of Ashlagi et al. [Ashlagi et al. 2019] this improves to $\widetilde{O}(\sqrt{n})$ given some additional constraints, and further improves to $O(\ln^4 n)$ in a tiered separable market [Ashlagi et al. 2019]. We note that for the bounded derivatives setting, with high probability, the communication cost will be $O(n^{1/3} \ln^{2/3} n)$ for all agents except the bottommost $\Theta(n^{2/3} \ln^{1/3} n)$, for whom the cost can reach $O(n^{2/3} \ln^{1/3} n)$.

Another approach to selecting which universities to apply to was considered by Shorrer who devised a dynamic program to compute the optimal choices for students assuming universities had a common ranking of students [Shorrer 2019].

2.2 Preliminaries

2.2.1 Some useful notation and definitions

There are n men and n women. In all of our models, each man m has a utility $U_{m,w}$ for the woman w, and each woman w has a utility $V_{m,w}$ for the man m. These utilities are defined as

$$U_{m,w} = U(r_w, s_m(w)), \text{ and}$$
$$V_{m,w} = V(r_m, s_w(m)),$$

where r_m and r_w are common public ratings, $s_m(w)$ and $s_w(m)$ are private scores specific to the pair (m, w), and $U(\cdot,\cdot)$ and $V(\cdot,\cdot)$ are continuous and strictly increasing functions from $\mathbb{R}_+^2$ to $\mathbb{R}_+$. The ratings are independent uniform draws from $[0, 1]$ as are the scores.

In the *Linear Separable Model*, each man m assigns each woman w a utility of $U_{m,w} = \lambda \cdot r_w +$

$(1-\lambda)\cdot s_m(w)$, where $0 < \lambda < 1$ is a constant. The women's utilities for the men are defined analogously as $V_{m,w} = \lambda \cdot r_m + (1-\lambda)\cdot s_w(m)$. All our experiments are for this model.

We let $\{m_1, m_2, \ldots, m_n\}$ be the men in descending order of their public ratings and $\{w_1, w_2, \ldots, w_n\}$ be a similar ordering of the women. We say that m_i has public rank i, or rank i for short, and similarly for w_i. We also say that m_i and w_i are *aligned*. In addition, we often want to identify the men or women in an interval of public ratings. Accordingly, we define $M(r, r')$ to be the set of men with public ratings in the range (r, r'), and $M[r, r']$ to be the set with public ratings in the range $[r, r']$; we also use the notation $M(r, r']$ and $M[r, r')$ to identify the men with ratings in the corresponding semi-open intervals. We use an analogous notation, with W replacing M, to refer to the corresponding sets of women.

We will be comparing the achieved utilities in stable matchings to the following benchmarks: the rank i man has as benchmark $U(r_{w_i}, 1)$, the utility he would obtain from the combination of the rank i woman's public rating and the highest possible private score; and similarly for the women. Based on this we define the loss an agent faces as follows.

Definition 1 (Loss). *Suppose man m and woman w both have rank i. The* loss *m sustains from a match of utility u is defined to be $U(r_w, 1) - u$. The loss for women is defined analogously.*

In our analysis we will consider a complete bipartite graph whose two sets of vertices correspond to the men and women, respectively. For each man m and woman w, we view the possible matched pair (m, w) as an edge in this graph. Thus, throughout this work, we will often refer to edges being proposed, as well as edges satisfying various conditions.

2.3 Upper Bound in The Linear Separable Model

To illustrate our proof technique for deriving upper bounds, we begin by stating and proving our upper bound result for the special case of the linear separable model with $\lambda = \frac{1}{2}$.

Theorem 2.11. *In the linear separable model with $\lambda = 1/2$, when there are n men and n women, for any given constant $c > 0$, for large enough n, with probability at least $1 - n^{-c}$, in every stable matching, for every i, with $r_{w_i} \geqslant \overline{\sigma} \triangleq 3\overline{L}/2$, agent m_i suffers a loss of at most $\overline{L}$, where $\overline{L} = (16(c + 2)\ln n/n)^{1/3}$, and similarly for the agents w_i.*

In words, w.h.p., all but the bottommost agents (those whose aligned agents have public rating less than $\overline{\sigma}$) suffer a loss of no more than $\overline{L}$. This is a special case of our basic upper bound for the bounded utilities model (Theorem 2.15).

One of our goals is to be able to limit the number of proposals the proposing side needs to make. We identify the edges that could be in some stable matching, calling them acceptable edges. Our definition is stated generally so that it covers all our results; accordingly we replace the terms $\overline{L}$ and $\overline{\sigma}$ in Theorem 2.11 with parameters L and σ.

Definition 2 (Acceptable edges)**.** *Let $0 < \sigma < 1$ and $0 < L < 1$ be two parameters. An edge (m_i, w_j) is (L, σ)-*man-acceptable *either if it provides m_i utility at least $U(r_{w_i}, 1) - L$, or if $m_i \in M[0, \sigma)$. The definition of (L, σ)-woman-acceptable is symmetric. Finally, (m_i, w_j) is (L, σ)-*acceptable *if it is both (L, σ)-man and (L, σ)-woman-acceptable.*

To prove our various results, we choose L and σ so that w.h.p. the edges in every stable matching are (L, σ)-*acceptable*. We call this high probability event $\mathcal{E}$. We will show that if $\mathcal{E}$ occurs, then running DA on the set of acceptable edges, or any superset of the acceptable edges obtained via loss thresholds, produces the same stable matching as running DA on the full set of edges.

Theorem 2.12. *If $\mathcal{E}$ occurs, then running woman-proposing DA with the edge set restricted to the acceptable edges or to any superset of the acceptable edges obtained via loss thresholds (including the full edge set) result in the same stable matching.*

The implication is that w.h.p. a woman can safely restrict her proposals to her acceptable edges, or to any overestimate of this set of edges obtained by her setting an upper bound on the loss she is willing to accept. There is a small probability— at most n^{-c}—that this may result in a less good outcome, which can happen only if $\mathcal{E}$ does not occur. Note that Theorem 2.12 applies to every utility model we consider. Then, w.h.p., every stable matching gives each woman w, whose aligned agent m has public rating $r_m \geqslant \overline{\sigma} = \Omega((\ln n/n)^{1/3})$, a partner with public rating in the range $[r_m - 2\overline{L}, r_m + \frac{5}{2}\overline{L}]$ (see Theorem 2.24 in Section 2.9.1). The bound $r_m - 2\overline{L}$ is a consequence of the bound on the woman's loss; the bound $r_m + \frac{5}{2}\overline{L}$ is a consequence of the bound on the men's losses. An analogous statement applies to the men.

This means that if we are running woman-proposing DA, each of these women might as well limit her proposals to her woman-acceptable edges, which is at most the men with public ratings in the range $r_m \pm \Theta(\overline{L})$ for whom she has private scores of at least $1 - \Theta(\overline{L})$. In expectation, this yields $\Theta(n^{1/3}(\ln n)^{2/3})$ men to whom it might be worth proposing. It also implies that a woman can have a gain of at most $\Theta(\overline{L})$ compared to her target utility.

If, in addition, each man can inexpensively signal the women who are man-acceptable to him, then the women can further limit their proposals to just those men providing them with a signal; this reduces the expected number of proposals these women can usefully make to just $\Theta(\ln n)$.

2.4 Proof of Theorem 2.11

We begin by outlining the main ideas used in our analysis. Our goal is to show that when we run woman proposing DA, w.h.p. each man receives a proposal that gives him a loss of at most L (except possibly for men among the bottommost $\Theta(nL)$). As the outcome is the man-pessimal stable matching, this means that w.h.p., in all stable matchings, these men have a loss of at most L. By symmetry, the same bound holds for the women.

Next, we provide some intuition for the proof of this result. See Fig. 2.1. Our analysis uses 3 parameters $\alpha, \beta, \gamma = \Theta(L)$. Let m_i be a non-bottommost man. We consider the set of men with public rank at least $r_{m_i} - \alpha$: $M_i = M[r_{m_i} - \alpha, 1]$. We consider a similar, slightly larger set of

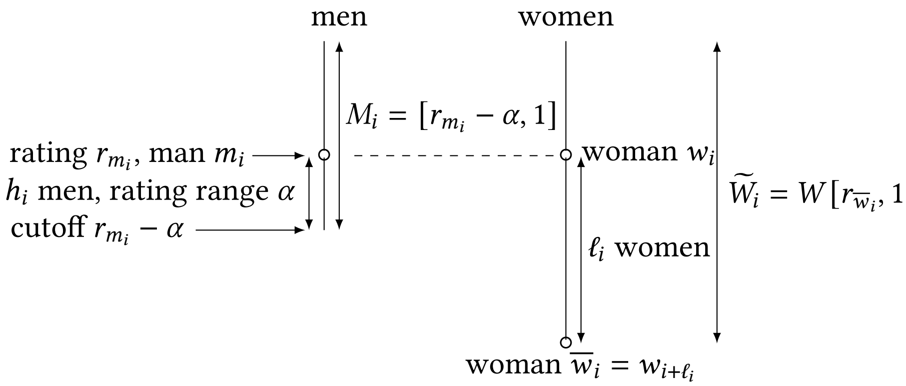

Figure 2.1: Illustrating Lemma 2.13

women: $\widetilde{W}_i = W[r_{w_i} - 3\alpha, 1]$. Now we look at the best proposals by the women in $\widetilde{W}_i$, i.e. the ones they make first. Specifically, we consider the proposals that give these women utility at least $V(r_{m_i} - \alpha, 1)$, proposals that are therefore guaranteed to be to the men in M_i. Let $|M_i| = i + h_i$ and $|W_i| = i + \ell_i$. In expectation, $\ell_i - h_i = 2\alpha n$. Necessarily, at least $\ell_i - h_i + 1$ women in M_i cannot match with men in $M_i \setminus \{m_i\}$. But, as we will see, these women all have probability at least β of having a proposal to m_i which gives them utility at least $V(r_{m_i} - \alpha, 1)$. These are proposals these women must make before they make any proposals to men with public rating less than $r_{m_i} - \alpha$. Furthermore, for each of these proposals, m_i has probability at least γ of having a loss of L or less. Thus, in expectation, m_i receives at least $2\alpha\beta\gamma n$ proposals which give him a loss of L or less.

We actually want a high-probability bound. So we choose α, β, γ so that $\alpha\beta\gamma n \geqslant c \log n$ for a suitable constant $c > 0$, and then apply a series of Chernoff bounds. There is one difficulty. The Chernoff bounds requires the various proposals to be independent. Unfortunately, in general, this does not appear to be the case. However, we are able to show that the failure probability for our setting is at most the failure probability in an artificial setting in which the events are independent, which yields the desired bound.

We now embark on the actual proof.

We formalize the men's rating cutoff with the notion of DA stopping at public rating r.

Definition 3 (DA stops)**.** *The women* stop at public rating r *if, in each woman's preference list, all the edges with utility less than* $V(r,1)$ *are removed. The women* stop at man m *if, in each woman's preference list, all the edges following her edge to* m *are removed. The women* double cut at man m and public rating r, *if they each stop at* m *or* r, *whichever comes first. Men stopping and double cutting are defined similarly. Finally, an edge is said to* survive *the cutoff if it is not removed by the stopping.*

To obtain our bounds for man m_i, we will have the women double cut at rating $r_{m_i} - \alpha$ and at man m_i, where $\alpha > 0$ is a parameter we will specify later.

Our upper bounds in all of the utility models depend on a parameterized key lemma (Lemma 2.13) stated shortly. This lemma concerns the losses the men face in the woman-proposing DA; a symmetric result applies to the women. The individual theorems follow by setting the parameters appropriately. Our key lemma uses three parameters: $\alpha, \beta, \gamma > 0$. To avoid rounding issues, we will choose α so that αn is an integer. The other parameters need to satisfy the following constraints.

$$\text{for } r \geqslant \alpha: \qquad V(r-\alpha, 1) \leqslant V(r, 1-\beta) \tag{2.1}$$

$$\text{for } r \geqslant 3\alpha: \qquad U(r,1) - U(r-3\alpha, 1-\gamma) \leqslant L \tag{2.2}$$

Equation (2.1) relates the range of private values that will yield a woman an edge to m_i that survives the cut at $r_{m_i} - \alpha$, or equivalently the probability of having such an edge. Observation 1 below, shows that Equation (2.2) identifies the range of m_i's private values for proposals from $\widetilde{W}_i$ that yield him a loss of at most L (for we will ensure the women in $\widetilde{W}_i$ have public rating at least $r_{w_i} - 3\alpha$).

Observation 1. *Consider the proposal from woman* w *to the rank* i *man* m_i. *Suppose the rank* i *woman* w_i *has rating* $r_{w_i} \geqslant 3\alpha$. *If* w *has public rating* $r \geqslant r_{w_i} - 3\alpha$ *and* m_i*'s private score for* w *is at least* $1-\gamma$, *then* m_i*'s utility for* w *is at least* $U(r_{w_i} - 3\alpha, 1-\gamma) \geqslant U(r_{w_i}, 1) - L$.

In the linear separable model with $\lambda = \frac{1}{2}$, we set $\alpha = \beta = \gamma$ and $L = 2\alpha$.

The next lemma determines the probability that man m_i receives a proposal causing him a loss of at most L. The lemma calculates this probability in terms of the parameters we just defined. Note that the result does not depend on the utility functions $U(\cdot,\cdot)$ and $V(\cdot,\cdot)$ being linear. In fact, the same lemma applies to much more general utility models which we also study (see Section 2.6) and it is the crucial tool we use in all our upper bound proofs.

In what follows, to avoid heavy-handed notation, by $r_{m_i} - \alpha$ we will mean $\max\{0, r_{m_i} - \alpha\}$.

In order to state our next result crisply, we define the following Event $\mathcal{E}_i$. It concerns a run of woman-proposing DA with double cut at the rank i man m_i and at public rating $r_{m_i} - \alpha$. Let $h_i = \left|M[r_{m_i} - \alpha, r_{m_i})\right|$, $\ell_i = \left|W[r_{w_i} - 3\alpha, r_{w_i})\right|$, and $\overline{w}_i$ be the woman with rank $i + \ell_i$. See Figure 2.1 for an illustration of these definitions. Event $\mathcal{E}_i$ occurs if $r_{w_i} \geqslant 3\alpha$ and between them the $i + \ell_i$ women in $W[r_{w_i} - 3\alpha, 1]$ make at least one proposal to m_i that causes him a loss of at most L.

Finally we define Event $\mathcal{E}$: it happens if $\mathcal{E}_i$ occurs for all i such that $r_{w_i} \geqslant 3\alpha$.

Lemma 2.13. *Let $\alpha > 0$ and $L > 0$ be given, and suppose that β and γ satisfy* (2.1) *and* (2.2)*, respectively. Then, Event $\mathcal{E}$ occurs with probability at least $1 - p_f$, where the failure probability*

$$p_f = n \cdot \exp(-\alpha(n-1)/12) + n \cdot \exp(-\alpha(n-1)/24) + n \exp(-\alpha\beta n/8) + n \cdot \exp(-\alpha\beta\gamma n/2).$$

The following simple claim notes that the men's loss when running the full DA is no larger than when running double-cut DA.

Claim 2.4.1. *Suppose a woman-proposing double-cut DA at man m_i and rating $r_{m_i} - \alpha$ is run, and suppose m_i incurs a loss of L. Then in the full run of woman-proposing DA, m_i will incur a loss of at most L.*

Proof. Recall that when running the women-proposing DA the order in which unmatched women are processed does not affect the outcome. Also note that as the run proceeds, whenever a man's match is updated, the man obtains an improved utility. Thus, in the run with the full edge set we can first use the edges used in the double-cut DA and then proceed with the remaining edges. Therefore if in the double-cut DA m_i has a loss of L, in the full run m_i will also have a loss of at most L. $\square$

To illustrate how this lemma is applied, we now prove Theorem 2.11. Note that $\overline{L}$ is the value of L used in this theorem. Our other results use other values of L.

Proof. (Of Theorem 2.11) By Lemma 2.13, in the double-cut DA, for all i with $r_{w_i} \geqslant 3\alpha$, m_i obtains a match giving him loss at most $\overline{L}$, with probability at least $1 - n \cdot \exp(-\alpha(n-1)/12) - n \cdot \exp(-\alpha n/24) - n \exp\left(-\alpha^2 n/8\right) - n \cdot \exp\left(-\alpha^3 n/2\right)$.

By Claim 2.4.1, m_i will incur a loss of at most $\overline{L}$ in the full run of woman-proposing DA with at least as large a probability. But this is the man-pessimal match. Consequently, in every stable match, m_i has a loss of at most $\overline{L}$. By symmetry, the same bound applies to each woman w_i such that $r_{m_i} \geqslant 3\alpha$.

We choose $\overline{L} = [16(c+2)\ln n/n]^{1/3}$. Recalling that $\alpha = \overline{L}/2$, we see that for large enough n the probability bound, over all the men and women, is at most $1 - n^{-c}$. The bounds $r_{w_i} \geqslant 3\alpha$ and $r_{m_i} \geqslant 3\alpha$ imply we can set $\overline{\sigma} = 3\alpha = \frac{3}{2}\overline{L}$. $\square$

Proof. (Of Lemma 2.13.) We run the double-cut DA in two phases, defined as follows. Recall that $h_i = \left|M[r_{m_i} - \alpha, r_{m_i})\right|$ and $\ell_i = \left|W[r_{w_i} - 3\alpha, r_{w_i})\right|$. Note that women with rank at most $i + \ell_i$ have public rating at least $r_{w_i} - 3\alpha$.

Phase 1. Every unmatched woman with rank at most $i+\ell_i$ keeps proposing until her next proposal is to m_i, or she runs out of proposals.

Phase 2. Each unmatched women makes her next proposal, if any, which will be a proposal to m_i.

Our analysis is based on the following four claims. The first two are simply observations that w.h.p. the number of agents with public ratings in a given interval is close to the expected number.

A critical issue in this analysis is to make sure the conditioning induced by the successive steps of the analysis does not affect the independence needed for subsequent steps. To achieve this, we use the Principle of Deferred Decisions, only instantiating random values as they are used. Since each successive bound uses a different collection of random variables this does not present a problem.

Claim 2.4.2. *Let $\mathcal{B}_1$ be the event that for some i, $h_i \geqslant \frac{3}{2}\alpha(n-1)$. $\mathcal{B}_1$ occurs with probability at most $n \cdot \exp(-\alpha(n-1)/12)$. The only randomness used in the proof are the choices of the men's public ratings. The same bound applies to the women.*

Proof. We prove the bound for an arbitrary man m with public rating r_m. The expected number n_x of men other than m in $M[r_m - \alpha, r_m]$ is $\alpha(n-1)$. This bound depends on the independent random choices of the men's public ratings. Thus, by a Chernoff bound,

$$\Pr\left[n_x \geqslant \tfrac{3}{2}\alpha(n-1)\right] \leqslant \exp(\alpha(n-1)/12).$$

Now, we apply a union bound to all n men to obtain the stated result. □

Claim 2.4.3. *Let $\mathcal{B}_2$ be the event that for some i, $\ell_i \leqslant \frac{5}{2}\alpha(n-1)$. $\mathcal{B}_2$ occurs with probability at most $n \cdot \exp(-\alpha(n-1)/24)$. The only randomness used in the proof are the choices of the women's public ratings. The same bound applies to the men.*

Proof. This proof is very similar to the proof of Claim 2.4.2. We prove the bound for an arbitrary woman w with public rating $r_w \geqslant 3\alpha$. The expected number n_y of women other than w in $W[r_w - 3\alpha, r_m]$ is $3\alpha(n-1)$. This bound depends on the independent random choices of the women's public ratings. Thus, by a Chernoff bound,

$$\Pr\left[n_y \leqslant \tfrac{5}{2}\alpha(n-1)\right] \leqslant \exp(\alpha(n-1)/24).$$

Now, we apply a union bound to all n women to obtain the stated result. $\square$

Claim 2.4.4. *Let $\mathcal{B}_3$ be the event that between them, the women with rank at most $i+\ell_i$ make fewer than $\frac{1}{2}\alpha\beta n$ Step 2 proposals to m_i. If events $\mathcal{B}_1$ and $\mathcal{B}_2$ do not occur, then $\mathcal{B}_3$ occurs with probability at most $\exp(-\alpha\beta n/8)$. The only randomness used in the proof are the choices of the women's private scores.*

This bound uses the private scores of the women and employs a novel argument given below to sidestep the conditioning among these proposals.

Claim 2.4.5. *If none of the events $\mathcal{B}_1$, $\mathcal{B}_2$, or $\mathcal{B}_3$ occur, then at least one of the Step 2 proposals to m_i will cause him a loss of at most L with probability at least $1-(1-\gamma)^{\alpha\beta n/2} \geqslant 1-\exp(-\alpha\beta\gamma n/2)$. The only randomness used in the proof are the choices of the men's private scores.*

Proof. Note that each Phase 2 proposal is from a woman w with rank at most $i+\ell_i$. As already observed, her public rating is at least $r_{w_i}-3\alpha$. Recall that man m_i's utility for w equals $U(r_w, s_{m_i}(w)) \geqslant U(r_{w_i}-3\alpha, s_{m_i}(w))$. To achieve utility at least $U(r_{w_i},1)-L \leqslant U(r_{w_i}-3\alpha, 1-\gamma)$ (using (2.2)) it suffices to have $s_{m_i}(w) \geqslant 1-\gamma$, which happens with probability γ. Consequently, utility at least $U(r_{w_i},1)-L$ is achieved with probability at least γ.

For each Phase 2 proposal these probabilities are independent as they reflect m_i's private scores for each of these proposals. Therefore the probability that there is no proposal providing m_i a loss of at most L is at most

$$\left(1-\gamma\right)^{\alpha\beta n/2} \leqslant \exp(\alpha\beta\gamma n/2).$$

$\square$

Concluding the proof of Lemma 2.13: The overall failure probability summed over all n choices of i is

$$n \cdot \exp(-\alpha(n-1)/12) + n \cdot \exp(-\alpha(n-1)/24) + n \exp(-\alpha\beta n/8) + n \cdot \exp(-\alpha\beta\gamma n/2).$$

□

Proof. (Of Claim 2.4.4.) First, we simplify the action space by viewing the decisions as being made on a discrete utility space, as specified in the next claim.

Claim 2.4.6. *For any $\delta > 0$, there is a discrete utility space in which for each woman the probability of selecting m_i is only increased, and the probability of having any differences in the sequence of actions in the original continuous setting and the discrete setting is at most δ.*

Proof. For each man m we partition the interval $[V(r_m, 0), V(r_m, 1)]$ of utilities it can provide into the following z subintervals: $[V(r_m, 0), V(r_m, 1/z)), [V(r_m, 1/z), V(r_m, 2/z)), \ldots, [V(r_m, (z-2)/z), V(r_m, (z-1)/z)), [V(r_m, (z-1)/z), V(r_m, 1)]$. Note that the probability that woman w's edge to m occurs in any one subinterval is $1/z$. Over all n men this specifies $n(z-1)$ utility values that are partitioning points. Now, for each man m, we partition the interval $[V(r_m, 0), V(r_m, 1)]$ about all $n(z-1)$ of these points, creating $n(z-1)+1$ subintervals. The values at these partition points plus the endpoint $V(r_m, 0)$ are the discrete utilities available to the women for evaluating man m, obtained by rounding down her actual utility.

Consider a single interval $I = [V(r_m, a), V(r_m, b))$ and an arbitrary woman w. Let $p_j^{I,c}$ be the probability that in the original continuous private score setting, the probability exactly one man m_j provides her a utility in I, let $p_{\text{none}}^{I,c}$ be the probability no one provides her a utility in I, and let $\overline{p}^{I,c}$ be the probability that two or more men provide her a utility in I. Note that $\overline{p}^{I,c} \leqslant n(n-1)/2z^2$. In the discrete setting, we remove the possibility of making two proposals and increase the probability of selecting man m_i by this amount: the probability of selecting man $m_j \neq m_i$ alone, with private score a will be $p_j^{I,c}$, the probability of selecting no one will be $p_{\text{none}}^{I,c}$, while the probability of selecting man m_i with private score a becomes $p_i^{I,d} = p_i^{I,c} + \overline{p}^{I,c}$.

Recall that in the run of double-cut DA, each woman repeatedly makes the next highest utility proposal. We view this as happening as follows. For each successive discrete utility value, woman w has the following choices.

i. she selects some man to propose to (among the men w she has not yet proposed to); or

ii. she takes "no action". This corresponds to w making no proposal achieving the current utility.

Every run of DA in the continuous setting that does not have a woman selecting two men over the course of a single utility interval will result in the identical run in the discrete setting in terms of the order in which each woman proposes to the men. Thus, the probability that in the discrete setting w's action in terms of who she selects and in what order differs from her actions in the continuous setting is at most $n^3/2z \triangleq \delta/n$ (because, in each possible computation, w makes at most nz choices, and for each choice the probability difference is at most $n^2/2z^2$). Furthermore, the probability of selecting man m_i is only increased. So over all n women, the probability of anything changing is at most δ. Clearly, δ can be made arbitrarily small. □

We represent the possible computations of the double-cut DA in this discrete setting using a tree T. Each woman will be going through her possible utility values in decreasing order, with the possible actions of the various women being interleaved in the order given by the DA processing. Each node u corresponds to a woman w processing her next utility value. The possible choices at this utility are each represented by an edge descending from u. These choices are:

i. Proposing to some man (among those men w has not yet proposed to); or

ii. "no action". This corresponds to w making no proposal achieving the current utility.

We observe the following important structural feature of tree T. Let S be the subtree descending from the edge corresponding to woman w proposing to m_i; in S there are no further actions of w, i.e. no nodes at which w makes a choice, because the double cut DA cuts at the proposal to m_i.

The assumption that $\mathcal{B}_1$ and $\mathcal{B}_2$ do not occur means that for all i, $h_i < \frac{3}{2}\alpha(n-1)$ and $\ell_i > \frac{5}{2}\alpha(n-1)$, and therefore $\ell_i - h_i > \alpha(n-1)$.

At each leaf of T, up to $i + h_i - 1$ women will have been matched with someone other than m_i. The other women either finished with a proposal to m_i or both failed to match and did not propose to m_i. Let w be a woman in the latter category. Then, on the path to this leaf, w will have traversed edges corresponding to a choice at each discrete utility in the range $[V(r_{m_i} - \alpha, 1), V(1, 1)]$.

We now create an extended tree, T_x, by adding a subtree at each leaf; this subtree will correspond to pretending there were no matches; the effect is that each women will take an action at all their remaining utility values in the range $[V(r_{m_i} - \alpha, 1), V(1, 1)]$, except that in the sub-subtrees descending from edges that correspond to some woman w selecting m_i, w has no further actions. For each leaf in the unextended tree, the probability of the path to that leaf is left unchanged. The probabilities of the paths in the extended tree are then calculated by multiplying the path probability in the unextended tree with the probabilities of each woman's choices in the extended portion of the tree.

Next, we create an artificial mechanism $\mathcal{M}$ that acts on tree T_x. The mechanism $\mathcal{M}$ is allowed to put $i + h_i - 1$ "blocks" on each path; blocks can be placed at internal nodes. A block names a woman w and corresponds to her matching (but we no longer think of the matches as corresponding to the outcome of the edge selection; they have no meaning beyond making all subsequent choices by this woman be the "no action" choice).

DA can be seen as choosing to place up to $i + h_i - 1$ blocks at each of the nodes corresponding to a leaf of T. $\mathcal{M}$ will place its blocks so as to minimize the probability p of paths with at least $\frac{1}{2}\alpha\beta n$ women choosing edges to m_i. Clearly p is a lower bound on the probability that the double-cut DA makes at least $\frac{1}{2}\alpha\beta n$ proposals in Step 2. Given a choice of blocks we call the resulting probability of having fewer than $\frac{1}{2}\alpha\beta n$ women choosing edges to m_i the *blocking probability.*

Claim 2.4.7. *The probability that $\mathcal{M}$ makes at least $\frac{1}{2}\alpha\beta n$ proposals to m_i is at least*
$1 - \exp(-\alpha\beta n/8)$.

Corollary 1. *The probability that the double-cut DA makes at least $\frac{1}{2}\alpha\beta n$ proposals to m_i is at least* $1 - \exp(-\alpha\beta n/8)$.

Proof. For any fixed δ, by Claim 2.4.7, the probability that $\mathcal{M}$ makes at least $\frac{1}{2}\alpha\beta n$ proposals to m_i is at least $1 - \exp(-\alpha\beta n/8)$. By construction, the probability is only larger for the double-cut DA in the discrete space.

Therefore, by Claim 2.4.1, the probability that the double-cut DA makes at least $\frac{1}{2}\alpha\beta n$ proposals to m_i in the actual continuous space is at least $1 - \exp(-\alpha\beta n/8) - \delta$, and this holds for any $\delta > 0$, however small. Consequently, this probability is at least $1 - \exp(-\alpha\beta n/8)$. □

Proof. (Of Claim 2.4.7.) We will show that the most effective blocking strategy is to block as many women as possible before they have made any choices. This leaves at least $(i + \ell_i) - (i - 1 + h_i) \geqslant 1 + \alpha(n-1) \geqslant \alpha n$ women unmatched. Then, as we argue next, each of these remaining at least αn women w has independent probability at least β that their proposal to m_i is cutoff-surviving. To be cutoff-surviving, it suffices that $V(r_{m_i}, s_w(m_i)) \geqslant V(r_{m_i} - \alpha, 1)$. But we know by (2.1) that $V(r_{m_i} - \alpha, 1) \leqslant V(r_{m_i}, 1 - \beta)$, and therefore it suffices that $s_w(m_i) \geqslant 1 - \beta$, which occurs with probability β.

Consequently, in expectation, there are at least $\alpha\beta n$ proposals to m_i, and therefore, by a Chernoff bound, at least $\frac{1}{2}\alpha\beta n$ proposals with probability at least $\exp(-\alpha\beta n/8)$.

We consider the actual blocking choices made by $\mathcal{M}$ and modify them bottom-up in a way that only reduces the probability of there being $\frac{1}{2}\alpha\beta n$ or more proposals to m_i.

Clearly, $\mathcal{M}$ can choose to block the same maximum number of women on every path as it never hurts to block more women (we allow the blocking of women who have already proposed to m_i even though it does not affect the number of proposals to m_i).

Consider a deepest block at some node u in the tree, and suppose b women are blocked at u. Let v be a sibling of u. As this is a deepest block, there will be no blocks at proper descendants of u, and furthermore as there are the same number of blocks on every path, v will also have b

blocked women.

Observe that if there is no blocking in a subtree, then the probability that a woman makes a proposal to m_i is independent of the outcomes for the other women. Therefore the correct blocking decision at node u is to block the b women with the highest probabilities of otherwise making a proposal to m_i, which we call their *proposing probabilities*; the same is true at each of its siblings v.

Let x be u's parent. Suppose the action at node x concerns woman $\widetilde{w}_x$. Note that the proposing probability for any woman $w \neq \widetilde{w}_x$ is the same at u and v because the remaining sequence of actions for woman w is the same at nodes u and v, and as they are independent of the actions of the other women, they yield the same probability of selecting m_i at some point.

We need to consider a number of cases.

Case 1. w is blocked at every child of x.

Then we could equally well block w at node x.

Case 2. At least one woman other than $\widetilde{w}_x$ is blocked at some child of x.

Each such blocked woman w has the same proposing probability at each child of x. Therefore by choosing to block the women with the highest proposing probabilities, we can ensure that at each node either $\widetilde{w}_x$ plus the same $b-1$ other women are blocked, or these $b-1$ woman plus the same additional woman $w' \neq \widetilde{w}_x$ are blocked. In any event, the blocking of the first $b-1$ women can be moved to x.

Case 2.1. $\widetilde{w}_x$ is not blocked at any child of x.

Then the remaining identical blocked woman at each child of x can be moved to x.

Case 2.2. $\widetilde{w}_x$ is blocked at some child of x but not at all the children of x.

Notice that we can avoid blocking $\widetilde{w}_x$ at the child u of x corresponding to selecting m_i, as the proposing probability for $\widetilde{w}_x$ after it has selected m_i is 0, so blocking any other women would be at least as good. Suppose that $w \neq \widetilde{w}_x$ is blocked at node u.

Let v be another child of x at which $\widetilde{w}_x$ is blocked. Necessarily, $p_{v,\widetilde{w}_x}$, the proposing probability for $\widetilde{w}_x$ at node v, is at least the proposing probability $p_{v,w}$ for w at node v (for otherwise w would be blocked at node v); also, $p_{v,w}$ equals the proposing probability for w at every child of x including u; in addition, $p_{v,\widetilde{w}_x}$ equals the proposing probability for $\widetilde{w}_x$ at every child of x other than u. It follows that w is blocked at u and $\widetilde{w}_x$ can be blocked at every other child of x. But then blocking $\widetilde{w}_x$ at x only reduces the proposing probability.

Thus in every case one should move the bottommost blocking decisions at a collection of sibling nodes to a single blocking decision at their parent. □

□

2.5 Making Fewer Proposals

We identify a sufficient set of edges that contains all stable matchings, and on which the DA algorithm produces the same outcome as when it runs on the full edge set.

Definition 4 (Viable edges)**.** *An edge (m, w) is* man-viable *if, according to m's preferences, w is at least as good as the woman he is matched to in the man-pessimal stable match. Woman-viable is defined symmetrically. An edge is* viable *if it is both man and woman-viable. E_v is the set of all viable edges.*

Lemma 2.14. *Running woman-proposing DA with the edge set restricted to E_v and with any superset obtained via loss thresholds, including the full edge set, results in the same stable matching.*

Proof. Suppose a new stable matching, S, now exists in the restricted edge set: it could not be present when using the full edge set, therefore there must be a blocking edge (m, w) in the full edge set. But neither m nor w would have removed this edge when forming their restricted edge set since for both of them it is better than an edge they did not remove (the edge they are matched with in S).

It follows that w.h.p. the set of stable matchings is the same when using E_v (or any super set of it generated by truncation with larger loss thresholds) and the whole set. Thus woman-proposing DA run on the restricted edge set will yield the same stable matching as on the full edge set.

□

Proof. (Of Theorem 2.12.) If $\mathcal{E}$ occurs, the set of acceptable edges contains all the viable edges. Furthermore, the acceptable edges are defined by means of loss thresholds. The result now follows from Lemma 2.14. □

For some of the very bottommost agents, almost all edges may be acceptable. However, in the bounded derivatives model, with slightly stronger constraints on the derivatives, we also show (see Section 2.11) the existence of an ε-Bayes-Nash equilibrium in which all but a bottom $\Theta((\ln n/n)^{1/3})$ fraction of agents use only $\Theta(\ln n)$ edges, and all agents propose using at most $\Theta(\ln^2 n)$ edges, with $\varepsilon = O(\ln n/n^{1/3})$.

2.6 More General Models

2.6.1 Utility Models

The General Utilities Model There are n men and n women. Each man m has a utility $U_{m,w}$ for the woman w, and each woman w has a utility $V_{m,w}$ for the man w. These utilities are defined as

$$U_{m,w} = U(r_w, s_m(w)), \text{ and}$$
$$V_{m,w} = V(r_m, s_w(m)),$$

where r_m and r_w are common public ratings, $s_m(w)$ and $s_w(m)$ are private scores specific to the pair (m, w), and $U(\cdot,\cdot)$ and $V(\cdot,\cdot)$ are continuous and strictly increasing functions from $\mathbb{R}_+^2$ to $\mathbb{R}_+$.

The public ratings and private scores are drawn independently from distributions with positive density functions with bounded support on $\mathbb{R}_+$. We assume without loss of generality that all public ratings and private scores are drawn uniformly and independently from $[0, 1]$ since there is always a change of variables that transforms them into uniform draws while transforming the utility functions monotonically.

U and V are not explicitly assumed to be bounded. However they are continuous and for the purpose of our analysis we can restrict the domain of U and V to the product of the bounded supports of our ratings and score distributions. These restricted U and V are continuous functions on a compact set and hence are bounded. Now, WLOG, by scaling appropriately, we can assume the range of U and V are both $[0, 1]$.

The Bounded Derivatives Model We add a notion of bounded derivatives to the general utilities model.

Definition 5. *A function $f(x, y) : \mathbb{R}^2 \to \mathbb{R}^+$ has (ρ, μ)-*bounded derivatives *if for all $(x, y) \in \mathbb{R}^2$,*

$$\begin{aligned} &\rho \leqslant \frac{\partial f}{\partial x} \Big/ \frac{\partial f}{\partial y} && \text{— the ratio bound;} \\ &\frac{\partial f}{\partial x} \leq \mu && \text{— the first derivative bound.} \end{aligned}$$

Note that this definition implies $\frac{\partial f}{\partial y}$ is upper bounded by μ/ρ.

In the bounded derivatives model, the utility functions U and V are restricted to having (ρ, μ)-bounded derivatives, for some constants $\rho, \mu > 0$. In the linear separable model, which is a special case of this model, $\mu = \lambda$ and $\rho = \lambda/(1 - \lambda)$.

2.6.2 Other Generalizations

Unequal numbers of men and women We generalize the above models to allow for n women and p men, where n and p need not be equal. Suppose that $n \leqslant p$. It is then convenient to change the public rating ranges to be $[0, p/n]$ for the men and $[p/n - 1, p/n]$ for the women. We proceed symmetrically when $n > p$. We will keep the private score range at $[0, 1]$. The effect of this change is to ensure that with high probability the top n public ratings for the men cover approximately the same range as the women's public ratings.

MANY-TO-ONE MATCHINGS The stable matching problem has also been studied in the setting of many-to-one matchings. For example, in the setting of employees and employers, often employers want to hire multiple employees. For this setting, we will refer to the two sides as companies and workers. Also, we will focus on the bounded derivatives setting.

There are n_c companies and n_w workers. Each company has d positions, meaning that it wants to match with d workers. Each worker can be hired by only one company. The total capacity of all the companies exactly matches the number of workers, i.e. $n_c \cdot d = n_w$.[3]

r_c will denote the public rating of company c, and r_w the public rating of worker w. Worker w has private score $s_w(c)$ for company c, and company c has private score $s_c(w)$ for worker w. $U(r_w, s_c(w))$ denotes the utility company c has for worker w, and $V(r_c, s_w(c))$ denotes the utility worker w has for company c.

To define the loss in the many-to-one setting, we need to define a non-symmetric notion of alignment of workers and companies.

Definition 6 (Alignment)**.** *Suppose company c has rank i (as per its public rating). Let w be the worker of rank $d \cdot i$ (also as per its public rating). Then w is aligned with c. Likewise, suppose worker w' has rank j. Let c' be the company with rank $\lceil j/d \rceil$. Then c' is aligned with w'.*

Definition 7 (Loss, cont.)**.** *Let c be a company and let w be aligned with c. The* loss *c sustains from a match of utility u is defined to be $U(r_w, 1) - u$. Similarly, let w' be a worker and let c' be aligned with w'. The* loss *w' sustains from a match of utility u is defined to be $V(r_{c'}, 1) - u$.*

[3]Our results generalize easily to the case in which the number of workers differs from the number of available positions. We omit the details.

2.7 Results

The Bounded Derivatives Model We begin by stating our basic result for this model.

Theorem 2.15. *In the bounded derivatives model, when there are n men and n women, for any given constant $c > 0$, for large enough n, with probability at least $1 - n^{-c}$, in every stable match, for every i, if $r_{w_i} \geqslant \overline{\sigma} \triangleq 3\overline{L}/4\mu$, agent m_i suffers a loss of at most $\overline{L}$, where $\overline{L} = \Theta((\ln n/n)^{1/3})$, and similarly for the agents w_i.*

Note that w.h.p., the public ratings of aligned agents are similar.

In words, w.h.p., all but the bottommost agents (those whose aligned agent has public rating less than $\overline{\sigma}$) suffer a loss of no more than $\overline{L}$. We call this high probability outcome $\mathcal{E}$.

By Theorem 2.12, the implication is that w.h.p. a woman can safely restrict her proposals to her acceptable edges, or to any overestimate of this set of edges obtained by her setting an upper bound on the loss she will accept from a match. There is a small probability— at most n^{-c}—that this may result in a less good outcome, namely the probability that $\mathcal{E}$ does not occur.

Then, w.h.p., every stable match gives each woman w, whose aligned agent m has public rating $r_m \geqslant \overline{\sigma} = \Omega((\ln n/n)^{1/3})$, a partner with public rating in the range $[r_m - \overline{L}/\mu, r_m + 5/4\overline{L}/\mu]$ (see Section 2.9.1). An analogous statement applies to the men.

This means that if we are running woman-proposing DA, each of these women might as well limit her proposals to her woman-acceptable edges, which is at most the men with public ratings in the range $r_m \pm \Theta(\overline{L})$ for whom she has private scores of at least $1 - \Theta(\overline{L})$. In expectation, this yields $\Theta(n^{1/3}(\ln n)^{2/3})$ men to whom it might be worth proposing. It also implies that a woman can have a gain of at most $\Theta(\overline{L})$ compared to her target utility.

If, in addition, each man can inexpensively signal the women who are man-acceptable to him, then the women can further limit their proposals to just those men providing them with a signal; in the case of accurate signals, this reduces the expected number of proposals these women can usefully make to just $\Theta(\ln n)$.

Our next result provides a distribution bound on the losses. It states that for most agents, the losses are at most $\Theta(\overline{L}/(\ln n)^{1/3})$, with a geometrically decreasing number of agents facing larger losses.

Theorem 2.16. *In the bounded derivatives model, when there are n men and n women, for any given constant $c > 0$, for large enough n, with probability at least $1-n^{-c}$, in every stable match, among the agents whose aligned partner has public score at least $\overline{\sigma} \triangleq 3\overline{L}/4\mu$, at most $2n \cdot \exp\big(-(c+2)\ln n/2^{3h}\big)$ men suffer a loss of more than $\overline{L}/2^h$, for integer h with $\frac{1}{3}\log\left(\frac{(c+2)\ln n}{\ln[n/3(c+2)\ln n]}\right) \leqslant h \leqslant \frac{1}{3}\log[(c+2)\ln n]$, and likewise for the women.*

We now generalize Theorem 2.15 to possibly unequal numbers of men and women, and also state what can be said for agents with low public ratings.

Theorem 2.17. *Suppose there are p men and w woman, with $p \geqslant n$. Let $t \geqslant 1$ be a parameter. In the bounded derivatives model, for any given constant $c > 0$, for large enough n, with probability at least $1 - n^{-c}$, in every stable match, every agent, except possibly the men whose aligned agents have public rating less than $\frac{p-n}{n} + \frac{\overline{\sigma}}{t}$ and the women whose aligned agents have public rating less than $\frac{\overline{\sigma}}{t}$, suffers a loss of at most $\overline{L}t^2$, where $\overline{L} = \Theta((\ln n/n)^{1/3})$ and $\overline{\sigma} = 3\overline{L}/4\mu$.*

Note that when $L = 1$ (i.e. 100% loss), $t = \Theta((n/\ln n)^{1/6})$ and therefore $\overline{\sigma}/t = \Theta((\ln n)/n)^{1/2})$, providing a lower bound on the range for which this result bounds the loss.

Setting $t = 1$ and $p = n$ yields Theorem 2.15.

The implication is similar to that for Theorem 2.15, but as t increases, i.e., for women whose aligned agents have increasingly low public ratings, the bound on the number of proposals she can usefully make grows by roughly a t^2 factor.

ε-Bayes-Nash Equilibrium

Definition 8. *A function* $f(x,y) : \mathbb{R}^2 \to \mathbb{R}^+$ *has* $(\rho_\ell, \rho_u, \mu_\ell, \mu_u)$-bounded derivatives *if for all* $(x,y) \in \mathbb{R}^2$,

$$\text{The ratio bound: } \rho_\ell \leqslant \frac{\partial f}{\partial x} \Big/ \frac{\partial f}{\partial y} \leqslant \rho_u.$$
$$\text{The first derivative bound: } \mu_\ell \leqslant \frac{\partial f}{\partial x} \leq \mu_u.$$

Then f *is said to have the* strong bounded derivative *property. Note that in the linearly separable model,* $\rho_\ell = \rho_u$ *and* $\mu_\ell = \mu_u$.

Let $t \geq 1$ be a parameter and $\overline{\sigma} = \Theta([\ln n/n]^{1/3})$. Define. $L_t^m \triangleq U(r_w, 1) - U(r_w - \overline{\sigma}t^2, 1)$ and $L_t^w \triangleq V(s_m, 1) - V(r_m - \overline{\sigma}t^2, 1)$. For this to be meaningful when $r_w - \overline{\sigma}t^2 < 0$, we extend the definition of U to this domain as follows. For $s < 0$, $\frac{\partial U(r,s)}{\partial r} = \mu_\ell$ and $\frac{\partial U(r,s)}{\partial s} = \rho_\ell$. We proceed analogously to handle the case that $r_m - \overline{\sigma}t^2 < 0$. Define parameters $\sigma_m = \beta/n^{1/3}$ and $\sigma_w = \nu/n^{1/3}$, where $\beta > 1$ and $\nu < 1$ are constants. We then define $t_m = \overline{\sigma}/\sigma_m$ and $t_w = \overline{\sigma}/\sigma_w$. Note that in the strongly bounded derivatives model, $L_{t_m}^m \leqslant \Theta\Big(\frac{\mu_u}{\beta^2} \cdot \frac{\ln n}{n^{1/3}}\Big)$ and $L_{t_w}^w \leqslant \Theta\Big(\frac{\mu_u}{\nu^2} \cdot \frac{\ln n}{n^{1/3}}\Big)$.

Theorem 2.18. *Let* $\varepsilon = \Theta(1/n^{1/3})$. *There are constants* $\beta > 1$ *and* $\nu < 1$ *such that in the strongly bounded derivatives model, there exists an* ε*-Bayes-Nash equilibrium where, with probability a least* $1 - n^c$, *agents with public ratings greater than* $\overline{\sigma}$ *make at most* $\Theta(\ln n)$ *proposals and all agents make at most* $\Theta(\ln^2 n)$ *proposals. Furthermore, in this equilibrium, with probability a least* $1 - n^c$, *every man has a loss of at most* $L_{t_m}^m$, *and every woman* w *has a loss of at most* $L_{t_w}^w$.

The General Utilities Model

Theorem 2.19. *Let* $0 < \varepsilon < 1$, $0 < \sigma < 1$, *and* $c > 0$ *be constants. In the general utilities model, for large enough n, with probability at least* $1 - \exp(-\Theta(n))$, *in every stable matching, every agent, except possibly those whose aligned agents have public rating less than* σ, *suffers a loss of at most* ϵ.

Clearly the smaller ε, the smaller the ranges of public ratings and private scores that can yield acceptable proposals; however, there does not appear to be a simple functional relationship between ε and the sizes of these ranges in this general model.

The Many to One Setting Next, we state our many-to-one result, expressing it in terms of the n_w workers and n_c companies, each having d positions. We now have possibly different bounds

$$\overline{L}_w = \Theta([(\ln n_w)/n_c]^{1/3}) = \Theta([(d \ln n_w)/n_w]^{1/3}), \quad \text{and}$$

$$\overline{L}_c = \begin{cases} \Theta((\max\{d, \ln n_w\}/n_w)^{1/3}) & d = O((n_w/\ln n_w)^{2/3}) \\ \Theta((d \ln n_w)/n_w) & d = \Omega(n_w/\ln n_w)^{2/3}) \end{cases}$$

on the losses for non-bottommost workers and companies. Analogous to the one-to-one case, we define $\overline{\sigma}_c = 3\overline{L}_c/4\mu$ and $\overline{\sigma}_w = 3\overline{L}_w/4\mu$, the public rating thresholds below which these loss bounds need not hold.

Theorem 2.20. *Let $\overline{L}_c$, $\overline{L}_w$, $\overline{\sigma}_c$ and $\overline{\sigma}_w$ be as defined above. Suppose that $d = O((n/\ln n)^{2/3})$. Then, for any given constant $k > 0$, with probability at least $1 - n^{-k}$, in every stable match, every company, except possibly those whose aligned agent has public rating less than $\overline{\sigma}_w$, suffers a loss of at most $\overline{L}_c$, and every worker, except possibly those whose aligned agent has public rating less than $\overline{\sigma}_c$, suffers a loss of at most $\overline{L}_w$.*

Lower Bounds The next two theorems show that the bounded derivative result is tight in two senses. First, we show that the bound $\overline{L}$ on the loss is tight up to a constant factor.

Theorem 2.21. *In the linear separable model with $\lambda = \frac{1}{2}$, if $n \geqslant 32,000$ and $L = \frac{1}{8}(\ln n/n)^{1/3}$, then with probability at least $\frac{1}{4}n^{-1/8}$ there is no perfect matching, let alone stable matching, in which every agent with public rating $\frac{3}{2}L$ or larger suffers a loss of at most L. (Here $\mu = \frac{1}{2}$, so $\frac{3}{2}L = 3L/4\mu$.)*

Next, we show that to obtain sub-constant losses in general, one needs constant bounds on the derivatives. We first define a notion of a sub-constant function, which we use to specify sub-constant losses.

Definition 9 (Sub-constant function). *A function $f(x) : \mathbb{R} \to \mathbb{R}^+$ is sub-constant if for every choice of constant $c > 0$, there exists an $\overline{x}$ such that for all $x \geqslant \overline{x}$, $f(x) \leqslant c$.*

Theorem 2.22. *Let $f : N \to \mathbb{R}^+$ be a continuous, strictly decreasing sub-constant function, and let $\delta, \sigma \in (0, 1)$ be constants. Then, in the following two cases, there exist continuous and strictly increasing utility functions $U(.,.)$ and $V(.,.)$ such that for some $\overline{n} > 0$, for all $n \geqslant \overline{n}$, with probability at least $1-\delta$, in every perfect matching, some rank i man m_i or woman w_i with public rating at least σ receives utility less than $U(r_{w_i}, 1) - f(n)$ or $V(r_{m_i}, 1) - f(n)$, respectively.*

i. $U(.,.)$ and $V(.,.)$ have derivatives w.r.t. their second variables that are bounded by a constant, but for (at least) one of which the derivative w.r.t. their first variable is not bounded by any constant.

ii. $U(.,.)$ and $V(.,.)$ have derivatives w.r.t. their first variables that are bounded by a constant, but for (at least) one of which the derivative w.r.t. their second variable is not bounded by any constant.

2.8 Proof Sketches for the Remaining Results

In Section 2.3, we proved Theorem 2.15 for the special case of the linear separable model with $\lambda = 1/2$. We will now briefly outline how we extend the analysis to the bounded derivative model and the general utilities model, as well as to the case where the number of men and women is unequal and the setting of many-to-one matchings. The full analyses can be found in Section 2.9.

We will also briefly discuss our construction of an ϵ-Bayes-Nash equilibrium in the bounded derivatives model as well as sketch our lower bound proofs in both the bounded derivatives and the general utility models. The full proofs can be found in Sections 2.11 and 2.10, respectively.

2.8.1 Extending the Upper Bound Result

1. Weaker bounds on the losses for agents with lower ranks.

This is obtained by reducing α to α/t, where $t > 1$, and replacing $\overline{L}$ by $L = 4\alpha t^2$. The only change occurs in recalculating the loss probability.

2. Unequal numbers of men and women.

The critical condition for the bound on m_i's loss is $r_{w_i} \geqslant 3\alpha$. This simply states that there is a range of 3α ratings below w_i. But this statement is independent of how many agents there are on each side. Similarly, the bound on w_i's loss requires that there be a range of 3α ratings below m_i. So all one has to do is rephrase these conditions in terms of p and n, the numbers of men and women, respectively.

3. The bounded derivatives model.

It suffices to scale the values of α, β, γ and L to take account of the bounded derivative property so as to ensure that Equations (2.1) and (2.2) still hold. As we shall see, setting $\beta = \alpha\rho$, $\gamma = \alpha\rho$ and $L = 4\alpha\mu$ suffices.

4. The many-to-one result.

We actually analyze the many-to-many setting. The main issue is that a company (replacing a man in the previous argument) seeks d_c matches rather than 1 and a worker seeks d_w matches. We need to restate Lemma 2.13, for now the alignment we seek is between positions sought by the workers and provided by the companies, rather than between men and women.

However, the significant change occurs in deducing the theorem, for now we need to determine the probability that a company receives d_c matches. The remaining changes are due to replacing n, the number of men and of women, with n_c and n_w, the numbers of companies and workers, respectively.

5. A distribution bound on the losses.

By reducing both α and L by a factor $s > 1$, we increase the failure probability for a single agent from $n^{-(c+1)}$ to $n^{-(c+1)/s^3}$. This implies, for example, that in expectation half the agents have a loss of $O(1/n^{1/3})$. In fact an analysis along the lines of observation (3) in the sketch proof shows that this bound holds with high probability.

2.8.2 Extensions to More General Models

1. The bounded derivative setting.

The only places we use the bounds on the derivatives are to determine β, γ, and L satisfying (2.1) and (2.2). As we show in Section 2.7, $\beta = \gamma = \alpha\rho$, and $L = 4\alpha/\mu$ suffice.

2. The general utilities model.

Given constants $\varepsilon, \sigma > 0$, we need to choose $\alpha, \beta, \gamma > 0$ and n large enough so that (2.1) and (2.2) are satisfied. The existence of such constant valued α, β, and γ follows using the fact that U and V, the utility functions, are continuous and strictly increasing.

2.8.3 Epsilon-Bayes-Nash Equilibrium

In the bounded derivative model, with slightly stronger constraints on the derivatives, we also show the existence of an ϵ-Bayes-Nash equilibrium in which agents make relatively few proposals. Specifically, there is an equilibrium in which no agent proposes more than $O(\ln^2 n)$ times and all but the bottommost $O((\ln n/n)^{1/3})$ fraction of the agents make only $O(\ln n)$ proposals. Here $\varepsilon = \Theta(\ln n/n^{1/3})$.

We use the idea of considering a run of DA with cuts just as in the proof of Theorem 2.15; in addition, the proposal receiving side will impose reservation thresholds based on their public rank. We also apply the distribution bound on losses described in (4) in the previous subsection. The resulting analysis is somewhat involved (see Section 2.11).

2.8.4 Lower Bounds

1. The lower bound complementing the one-to-one upper bound.

The main idea is to show by a direct computation that for each woman, with probability at least $1/n^{1/8}$, all her incident edges provide a loss of more than L to either her or her partner. We will need to exclude some low-probability events in which the number of agents in an interval is far from its expectation, and also eliminate the agents with public ratings less than $\frac{3}{2}L$. The net effect is that with probability at least $1/4n^{1/8}$ some woman has no incident $(L, \frac{3}{2}L)$-acceptable edge, where $L = \frac{1}{8}(\ln n/n)^{1/3}$, and hence with this probability there is no matching using solely $(L, \frac{3}{2}L)$-acceptable edges. Consequently, in order to obtain a stable matching with high probability, we need to increase the value of L.

2. The lower bound complementing the general utilities model upper bound.

To show that no sub-constant loss bound (such as $(\ln n/n)^{1/3}$) is possible, we consider a loss bound that is shrinking (slowly) as a function of n. For a given n, this can be expressed as a loss bound $\varepsilon(n)$. We provide two similar constructions as there are two separate derivative bounds.

Our first construction uses a utility function $U(s, v) = \frac{1}{2}(s + g(v))$, with $g(1) = 1$ and $g(\cdot)$ being unboundedly rapidly growing as $v \to 1$. g is designed to ensure that with high probability the edges to the women with public ratings $s \geqslant 1 - 2\varepsilon(n)$ all have private scores less than $1 - \varepsilon(n)$. This will ensure that with high probability m_1, the man with the highest public ranking, will have no edge providing him a loss of at most $\varepsilon(n)$. However slowly $\varepsilon(n)$ decreases as a function of n, we show that we can construct a corresponding g that grows suitably quickly. This construction demonstrates that the parameter ϵ needs to be constant. Notice that our construction actually shows that, in the general setting, w.h.p, there is not only no stable matching where all high public rating agents face sub-constant losses, but in fact no perfect matching.

2.9 Proofs of the Remaining Upper Bound Results

Proof. (Of Theorem 2.15.) We now consider what changes occur when we are no longer restricted to the linear separable model with $\lambda = \frac{1}{2}$.

First, we need to determine the values for β, γ and $\overline{L}$ implied by the bounded derivative parameters ρ and μ. We show the following values, $\beta = \alpha\rho$, $\gamma = \alpha\rho$ and $L = 4\alpha\mu$, satisfy (2.1) and (2.2).

For by the definition of ρ, $V(r-\alpha, 1) \leqslant V(r, 1-\alpha\rho) = V(r, 1-\beta)$, satisfying (2.1). And by the definition of ρ and μ, $U(r,1)-U(r-3\alpha, 1-\gamma) \leqslant U(r,1)-U(r-3\alpha-\gamma/\rho, 1) \leqslant (3\alpha+\gamma/\rho)\mu = 4\alpha\mu = L$, satisfying (2.2).

To complete the argument, it suffices to determine the failure probability on setting $L = \overline{L}$ when running the double-cut DA. Recall that the failure probability (summed over the $2n$ men and women) is given by:

$$\begin{aligned}
p_f &= 2n \cdot \exp(-\alpha(n-1)/12) + 2n \cdot \exp(-\alpha(n-1)/24) + 2n \exp(-\alpha\beta n/8) + 2n \cdot \exp(-\alpha\beta\gamma n/2) \\
&\leqslant 2n \cdot \exp(-\alpha(n-1)/12) + 2n \cdot \exp(-\alpha(n-1)/24) + 2n \exp\left(-\alpha^2\rho n/8\right) + 2n \cdot \exp\left(-\alpha^3\rho^2 n/2\right)
\end{aligned}$$

We note that $\alpha = \overline{L}/4\mu$, and set $\overline{L} = [128(c+2)\mu^3 \ln n/(\rho^2 n)]^{1/3}$. For large enough n this ensures a failure probability of at most n^{-c}. □

Proof. (Of Theorem 2.17.) We now need to consider smaller intervals of men and women below m_i and w_i respectively.

We set $\sigma = \overline{\sigma}/t$, where $t \geqslant 1$. We then set $\alpha = \sigma/4$ and $\beta = \alpha\rho$ as before, but to keep the most significant term in the probability bound unchanged ($2n \cdot \exp(-\alpha\beta\gamma n/2)$), we increase γ by a factor of t^2. We also set $L = \overline{L}t^2$.

The failure probability continues to be at most n^{-c} for large enough n so long as $\alpha\beta n = \Omega(c \ln n)$; this holds for $\sigma = \Omega((\ln n)/n)^{1/2}$.

Now let's consider what happens when there are p men and n women. We start with the case $p \geqslant n$. Our key lemma is stated w.r.t. the rank i man m_i and the rank i woman w_i, and requires $r_{w_i} \geqslant 3\alpha$ when the bottom of the rating range is 0 for both men and women.

It is convenient to have the range of ratings for the men be $[0, p/n]$ and for the women be $[p/n-1, p/n]$. The effect is that the expected values for r_{w_i} and r_{m_i} are equal. The condition for m_i to have a loss of at most L becomes $r_{w_i} \geqslant \frac{p}{n} + 3\alpha$ (i.e. w_i has a rating at least 3α greater than the bottommost possible rating for the women). But the condition for w_i to have a loss of at most L remains $r_{m_i} \geqslant 3\alpha$ (i.e. m_i has a rating at least 3α greater than the bottommost possible rating for the men).

Symmetric bounds apply when $n \geqslant p$. $\square$

Proof. (Of Theorem 2.19) We set $\alpha = \sigma/3$ and $L = \varepsilon$. Again, we need to satisfy (2.1) and (2.2).

To define β we begin by specifying a parameter $\beta(r,\alpha)$. There are two cases. If $V(r-\alpha,1) \leqslant V(r,0)$, then $\beta(r,\alpha) = 1$. Otherwise, as V is continuous and strictly increasing, there must be a value $\beta(r,\alpha) > 0$ such that $V(r-\alpha,1) = V(r,1-\beta(r,\alpha))$. Now, we define $\beta = \min_{r\in[\alpha,1]}\{\beta(r,\alpha)\}$. As this is the minimum of strictly positive values on a compact set, it follows that $\beta > 0$, also. Note that $V(r-\alpha,1) \leqslant V(r,1-\beta)$ for all $r \in [\alpha,1]$, satisfying (2.1). Also, $\beta = \Theta(1)$ if $\alpha = \Theta(1)$.

Similarly, to define γ we begin by specifying a parameter $\gamma(r,\alpha,\varepsilon)$. Again, there are two cases. If $U(r,1) - U(r-3\alpha,0) \leqslant L = \varepsilon$ then $\gamma(r,\alpha,\varepsilon) = 1$. Otherwise, as U is continuous and strictly increasing, there must be a value $\gamma(r,\alpha,\varepsilon) > 0$ such that $U(r,1) - U(r-3\alpha, 1-\gamma(r,\alpha,\varepsilon)) = \varepsilon$. Now, we define $\gamma \triangleq \min_{r\in[3\alpha,1]}\{\gamma(r,\alpha,\varepsilon)\}$. Again, as this is a minimum of strictly positive values on a compact set, $\gamma > 0$ also. Note that $U(r,1) - \varepsilon \leqslant U(r-3\alpha, 1-\gamma)$ for all $r \in [3\alpha,1]$, satisfying (2.2). Also, $\gamma = \Theta(1)$ if $\alpha = \Theta(1)$.

As $\sigma = \Theta(1)$, all of $\alpha,\beta,\gamma = \Theta(1)$. Therefore, by Lemma 2.13, the failure probability is $\exp(-\Theta(n))$. $\square$

Because the many-to-one setting is non-symmetric it is actually simpler to analyze the many-to-many setting, many-to-one being just a special case of this. We will use the terminology of workers and companies, for want of a better alternative. (One could think of these workers as being consultants or gig workers who seek multiple tasks at a time.)

In this setting there are n_c companies $c_1, c_2, \ldots, c_{n_c}$, and n_w workers, $w_1, w_2, \ldots, w_{n_w}$, both ordered by their public ranks. Each company has d_c tasks, and each worker desires d_w tasks. For simplicity, we suppose $n_c \cdot d_c = n_w \cdot d_w \triangleq n$. We let $n_{\max} = \max\{n_c, n_w\}$. There will be two loss parameters, $\overline{L}_c$, for the companies, and $\overline{L}_w$, for the workers. Finally, we use the notation $C(I)$ and $W(I)$, where I is an interval of public ratings, to denote, respectively, the companies and workers with public ratings in the interval I.

Definition 10 (Alignment)**.** *Suppose company c has rank i. Let w be the worker with rank $\lceil d_c \cdot i/d_w \rceil$. Then w is aligned with c. Likewise, suppose worker w' has rank j. Let c' be the company with rank $\lceil d_w \cdot j/d_c \rceil$. Then c' is aligned with w'.*

Definition 11 (company-acceptable edges)**.** *Let $0 < \sigma_c, \sigma_w < 1$, $0 < \overline{L}_c, \overline{L}_w < 1$ be parameters. An edge (c, w) is company-acceptable if either $c \in C[0, \sigma_c)$, or the utility c gets from this match is at least $U(r_{w'}, 1) - \overline{L}_c$, where w' is the worker aligned with c. Worker-acceptability requires either $c \in [0, \sigma_w)$, or utility at least $V(r_{c'}, 1) - \overline{L}_w$, where c' is the company aligned with w. An edge is acceptable if it is both company and worker-acceptable. (Strictly speaking, the definition is w.r.t. the four parameters σ_c, σ_w, $\overline{L}_c$, and $\overline{L}_w$, but for the sake of readability, we omit them from the terms company- and worker-acceptable.)*

Definition 12 (DA stops)**.** *The workers* stop at public rating *r if in each worker's preference list all the edges with utility less than $V(r, 1)$ are removed. The workers* stop at company *c if in each worker's preference list all the edges following their edge to c are removed. The workers* double cut at *c* and public rating *r, if they each stop at c or r, whichever comes first. Companies stopping and double cutting are defined similarly.*

Theorem 2.23. *Suppose that $d_w/d_c = O((n_w/\ln n_{\max})^{2/3})$. Then, in the bounded derivatives model, for any given constant $k > 0$, with probability at least $1 - n^{-k}$, in every stable matching, every company c_i, for which the aligned worker w_{j_i} has public rating at least $\sigma_w = \Theta(\overline{L}_w)$, suffers loss at most*

$$\overline{L}_c = \begin{cases} \Theta([(\ln n_{\max})/n_w]^{1/3}) & d_c = O(\ln n_{\max}) \\ \Theta([d_c/n_w]^{1/3}) & \ln n_{\max} \leqslant d_c = O(d_w(n_w/\ln n_{\max})^{2/3}) \\ \Theta(d_c \ln n_{\max}/[d_w n_w]) & d_c = \Omega(d_w(n_w/\ln n_{\max})^{2/3}) \end{cases}$$

and a corresponding symmetric bound for the workers' loss.

Proof. We need to take account of the fact that each company seeks to fill d_c positions and each worker seeks d_w positions. So we slightly redefine the double-cut DA to state that each worker who is not fully matched, i.e., who has fewer than d_w matches, keeps trying to match, stopping when she runs out of proposals, or she is fully matched, or her next proposal is to c_i.

First, to avoid rounding issues, we assume α is chosen so that $\alpha(n_w - \frac{5}{2})$ is an integer for the argument bounding $\overline{L}_c$, and similarly $\alpha(n_c - \frac{5}{2})$ is an integer for the argument bounding $\overline{L}_w$.

We introduce one more index: $j_i = \lceil d_c \cdot i/d_w \rceil$. We then define $\ell_i = \lfloor d_c(i + h_i - 1)/d_w \rfloor + \alpha(n_w - \frac{5}{2}) - j_i$ (this is where we use the assumption that $\alpha(n_w - \frac{5}{2})$ is an integer as ℓ_i has to be an integer). This will ensure that after running the double cut DA, the number of not fully matched workers is at least $\alpha(n_w - \frac{5}{2})$. To see this, note that the number of available positions is $d_c(i + h_i - 1)$ (remember c_i is not matched in Step 1); therefore, the number of fully matched workers is at most $\lfloor d_c(i+h_i-1)/d_w \rfloor$ and therefore the number of not-fully matched workers is at least $\ell_i + j_i - \lfloor d_c(i+h_i-1)/d_w \rfloor \geqslant \lfloor d_c(i+h_i-1)/d_w \rfloor + \alpha(n_w - \frac{5}{2}) - \lfloor d_c(i+h_i-1)/d_w \rfloor \geqslant \alpha(n_w - \frac{5}{2})$.

We need to make small changes to Claims 2.4.2–2.4.5 and to their proofs. It seems simplest to restate and, as necessary, reprove the claims.

Claim 2.9.1. *Let $\mathcal{B}_1$ be the event that for some i, $h_i = \left|C[r_{c_i} - \alpha, r_{c_i})\right| \geqslant \frac{3}{2}\alpha(n_c - 1)$. $\mathcal{B}_1$ occurs with* *probability at most* $n_c \cdot \exp(-\alpha(n_c - 1)/12)$*. The only randomness used in the proof are the choices of the companies' public ratings. An analogous bound applies to the workers.*

Its proof is unchanged. We just replace n by n_c.

Claim 2.9.2. *Let* $\mathcal{B}_2$ *be the event that for some* i, $\ell_i = \left|W[r_{w_i} - 3\alpha, r_{w_i})\right| \leqslant \frac{5}{2}\alpha(n_w - 1)$*. Then* $\mathcal{B}_2$ *occurs with probability at most* $n_w \cdot \exp(-\alpha(n_w - 1)/24)$*. The only randomness used in the proof are the choices of the workers' public ratings. An analogous bound applies to the companies.*

Its proof is unchanged. We just replace n by n_w.

Claim 2.9.3. *Let* $\mathcal{B}_3$ *be the event that between them, the workers with rank at most* $j_i + \ell_i$ *make at least* $\frac{1}{2}\alpha\beta(n_w - 5/2)$ *Step 2 proposals to* c_i*. If events* $\mathcal{B}_1$ *and* $\mathcal{B}_2$ *do not occur, then* $\mathcal{B}_3$ *occurs with probability at most* $\exp(-\alpha\beta(n_w - 5/2)/8)$*.*

It's proof is largely unchanged. The first issue is that now in the run of the DA algorithm placing a block on a worker w corresponds to w having matched d_w times. The proof is otherwise unchanged as any unblocked worker will run through her full utility range as before. However, the calculations change as follows. The number of not-fully matched workers is at least

$$\begin{aligned} d_w(j_i + \ell_i) - d_c(i - 1 + h_i) &\geqslant d_w \cdot \frac{5}{2}\alpha(n_w - 1) + d_c - d_c \cdot \frac{3}{2}\alpha(n_c - 1) \\ &\geqslant \alpha d_w n_w - \frac{5}{2}\alpha d_w \geqslant \alpha d_w(n_w - \frac{5}{2}). \end{aligned}$$

This causes the replacement of n by $n_w - 5/2$ in the bounds.

Claim 2.9.4. *If none of the events* $\mathcal{B}_1$, $\mathcal{B}_2$, *or* $\mathcal{B}_3$ *occur, then at least* $\frac{1}{4}\alpha\beta\gamma(n_w - 5/2)$ *of the Step 2 proposals to* c_i *will each cause* c_i *a loss of at most* $\overline{L}_c$ *with probability at least*
$1 - \exp(-\alpha\beta\gamma(n_w - 5/2)/16)$.

To ensure c_i receives at least d_c proposals that each cause it a loss at most $\overline{L}_c$, by Claim 2.9.4, we need that

$$\tfrac{1}{4}\alpha\beta\gamma\big(n_w - \tfrac{5}{2}\big) \geqslant d_c. \tag{2.3}$$

Then the overall failure probability summed over all companies is at most

$$n_c \cdot \exp(-\alpha(n_c - 1)/12) + n_w \cdot \exp(-\alpha(n_w - 1)/24) + n_c \exp(-\alpha\beta(n_w - 5/2)/8)$$
$$+ n_c \exp(-\alpha\beta\gamma(n_w - 5/2)/16).$$

In the bounded derivative setting, we continue to set $\beta = \alpha\rho$, $\gamma = \alpha\rho$ and $\overline{L}_c = 4\alpha\mu$. Then, for large enough n_c, n_w, with $\overline{L}_c \geqslant 4\mu \cdot [16(k+2)\ln n_{\max}/\rho^2(n_w - 5/2)]^{1/3}$ and $\overline{L}_c \geqslant 48\mu(k+2)\ln n_{\max}/(n_c - 1) \geqslant 47\mu(d_c/d_w)[(k+2)\ln n_{\max}/(n_w - 1)]$, the overall failure probability is at most $n_{\max}^{-k}$. The first of the two bounds on $\overline{L}_c$ dominates if $d_c/d_w = O((n_w/\ln n_{\max})^{2/3})$. In addition, with $\overline{L}_c \geqslant 4\mu \cdot (4d_c/\rho^2(n_w - \frac{5}{2}))^{1/3}$, (2.3) is satisfied. Thus, the overall condition is that $\overline{L}_c = \Omega(\max\{d_c, \ln n_{\max}\}/n_w)^{1/3})$.

The corresponding bound $\overline{L}_w = \Omega(\max\{d_w, \ln n_{\max}\}/n_c)^{1/3})$ can be deduced using the company-proposing DA.

It remains to prove Claim 2.9.4, which we do below. □

Proof. (Of Claim 2.9.4.) As $\mathcal{B}_3$ does not occur, by Claim 2.9.3, there are at least $\frac{1}{2}\alpha\beta(n_w - 5/2)$ Step 2 proposals to c_i. As explained in observation (4) of the sketch proof, each Step 2 proposal has independent probability at least γ of causing c_i a loss of at most $\overline{L}_c$ (the independence is because this is due to the private score of c_i for this proposal). In expectation, there are at least $\frac{1}{2}\alpha\beta\gamma(n_w - 5/2)$ of the proposals causing c_i a loss of at most $\overline{L}_c$, and by a Chernoff bound at least $\frac{1}{4}\alpha\beta\gamma(n_w - 5/2)$ such proposals to c_i with failure probability at most $\exp(-\alpha\beta\gamma(n_w - 5/2)/16)$. □

2.9.1 Range of Public Ratings for Acceptable Edges

Here we prove that in the one-to-one bounded derivative setting, with high probability, for each i, the acceptable edges from women w_i are to men with public rating in the range $[r_{m_i}-4\alpha, r_{m_i}+5\alpha]$, which we call w_i's *cone*. A symmetric bound applies to the men.

We then obtain a similar bound for the many-to-one setting.

Theorem 2.24. *In the one-to-one bounded derivative setting with n men and n women, for large enough n, with probability $1-n^{-c}$, for each i, the acceptable edges from women w_i are to men with public rating in the range $[r_{m_i}-4\alpha, r_{m_i}+5\alpha]$, where $\alpha = \overline{L}/4\mu$. A symmetric bound applies to the men.*

Proof. Theorem 2.15 bounds the loss by $\overline{L}$ for non-bottommost agents with probability $1-n^{-c}$ for large enough n. Therefore w_i will not be interested in matching with any man with public rating less than $r_{m_i} - \overline{L}/\mu$ (for any such man would give a loss greater than $\overline{L}$).

The situation to higher rated men needs a little more calculation. Let m_g be such a man. Then what matters is whether m_g incurs a loss of more than $\overline{L}$ if matched to w_i. This happens if $r_{w_g} - r_{w_i} > \overline{L}/\mu = 4\alpha$. We now show that to obtain $r_{w_g} - r_{w_i} \leqslant 4\alpha$, w.h.p. we must have $r_{m_g} - r_{m_i} \leqslant 5\alpha$.

We prove this in two steps: first, we show that w.h.p., if $r_{w_g} - r_{w_i} \leqslant 4\alpha$ then $i-g < 4\alpha(n-1) + \frac{1}{2}\alpha(n-1)$. Second, we show that w.h.p., $r_{m_g} - r_{m_i} \leqslant (4\alpha + \frac{1}{2}\alpha) + \frac{1}{2}\alpha = 5\alpha = \frac{5}{4}L/\mu$.

The expected number of women in $W[r_{w_g}-4\alpha, r_{w_g}]$ other than w_g is at most $4\alpha(n-1)$; and so by a Chernoff bound this number is at least $4\alpha(n-1)+\frac{1}{2}\alpha(n-1)$ with probability at most $\exp(-\alpha(n-1)/48)$. Call this bad event $\mathcal{B}_4$. Note that by assumption, $w_i \in W[r_{w_g}-4\alpha, r_{w_g}]$, and so if $\mathcal{B}_4$ does not occur $i-g < 4\frac{1}{2}\alpha(n-1)$.

Now suppose $\mathcal{B}_4$ does not occur, and consider the set $M[r_{m_g}-5\alpha, r_{m_g}]$. In expectation, other than m_g, it contains $5\alpha(n-1)$ men. By a Chernoff bound, it contains at most $4\frac{1}{2}\alpha(n-1)$ men other than m_g with probability at most $\exp(-\alpha(n-1)/40)$. Call this bad event $\mathcal{B}_5$.

If neither $\mathcal{B}_4$ nor $\mathcal{B}_5$ occur, as $i - g < 4\frac{1}{2}\alpha(n-1)$, $m_i \in M(r_{m_g} - 5\alpha, r_{m_g}]$, and therefore $r_{m_g} - r_{m_i} < 5\alpha \leqslant \frac{5}{4}\overline{L}/\mu$.

A union bound over the n women and n men gives a failure probability of $2n \cdot \exp(-\alpha(n-1)/48) + 2n \cdot \exp(-\alpha(n-1)/40)$, plus the failure probability from the proof of Theorem 2.15, which was actually at most

$$2n \cdot \exp(-\alpha(n-1)/12) + 2n \cdot \exp(-\alpha n/24) + 2n \exp\left(-\alpha^2 \rho n/8\right) + 2n \cdot \exp\left(-\alpha^3 \rho^2 n/2\right).$$

Even adding the new terms, for large enough n it still suffices to set $\overline{L} = [128(c{+}2)\mu^3 \ln n/(\rho^2 n)]^{1/3}$ to achieve an overall n^{-c} failure probability. □

We now extend the result to the many-to-many setting.

Theorem 2.25. *In the many-to-many bounded derivative setting with n_c companies and n_w workers, for large enough n_c and n_w, with probability $1 - n^{-k}$, for each i, the acceptable edges from worker w_i are to companies with public rating in the range $[r_{c_{j_i}} - L_w/\mu, r_{c_{j_i}} + 5L_c n_c/4\mu(n_c - 1)]$, where c_{j_i} is the company aligned with w_i. A symmetric bound applies to the companies.*

Proof. We need to adapt the previous proof to account for the fact that there are n_c companies and n_w workers.

The argument demonstrating the lower limit is unchanged, except we replace $\overline{L}$ with L_w. For the upper limit, we change the argument as follows. Now, we replace $\overline{L}$ by L_c .

We first observe that the number of workers in $W[r_{w_g} - 4\alpha, r_{w_g}]$ other than w_g is at least $4\alpha(n_w - 1) + \frac{1}{2}\alpha(n_w - 1)$ with probability at most $\exp(-\alpha(n_w - 1)/48)$. Call this bad event $\mathcal{B}_4$.

Second, if $\mathcal{B}_4$ does not occur, the set $M[r_{c_g} - 5\alpha n_c/(n_c - 1), r_{c_g}]$ contains at most $4\frac{1}{2}\alpha n_c$ companies other than c_g with probability at most $\exp(-\alpha n_c/40)$. Call this bad event $\mathcal{B}_5$.

Suppose neither $\mathcal{B}_4$ nor $\mathcal{B}_5$ occur. Then, the number of positions sought by the workers in $M[r_{w_{i-1}}, r_{w_g}]$ is $d_w(i-g) < 4\frac{1}{2}d_w\alpha(n_w-1) = 4\frac{1}{2}\alpha d_c n_c - 4\frac{1}{2}d_w\alpha$, while the number of positions available in $C[r_{c_g} - 5\alpha, r_{c_g}]$ is more than $4\frac{1}{2}d_c\alpha n_c = 4\frac{1}{2}\alpha d_c n_c$. Thus the number of available positions is at least the number sought, and therefore $c_i \in C[r_{c_g} - 5\alpha n_c/(n_c - 1), r_{c_g}]$.

It remains to revisit the probability bounds. The failure probability from the proof of Theorem 2.20 summed over all workers is

$$
\begin{aligned}
& n_w \cdot \exp(-\alpha(n_w - 1)/12) + \min\{n_c, n_w\} \cdot \exp(-\alpha(n_c - 1)/24) \\
& + n_w \exp(-\alpha\beta(n_c - 5/2)/8) + n_w \exp(-\alpha\beta\gamma(n_c - 5/2)/16),
\end{aligned}
$$

where $\beta = \gamma = \alpha\rho$. There is an analogous bound for the companies. Again, for large enough n_c and n_w, we can use the same values for L_c and L_w as before while maintaining the total failure probability at n^{-k}. □

2.9.2 Distribution Bound on Losses (Proof of Theorem 2.16)

Recall event $\mathcal{B}_3$ from the proof of Lemma 2.13, that the women with rank at most $i + \ell_i$ make fewer than $\frac{1}{2}\alpha\beta n$ Step 2 proposals to m_i. Claim 2.4.4 shows that if $\mathcal{B}_1$ and $\mathcal{B}_2$ do not occur then $\mathcal{B}_3$ occurs with probability at most $\exp(-\alpha\beta n/8)$.

Now, consider a man m and the aligned woman w, where $r_w \geqslant 4\alpha$. Let $\beta = \gamma = \alpha\rho$. We will bound the probability that m has a loss of more than $L^m_\alpha \triangleq U(r_w, 1) - (r_w - 4\alpha, 1)$.

If, in addition, $r_w \geqslant \overline{\sigma}$ and $\alpha = \overline{\sigma}/(4 \cdot 2^h) = \overline{L}/(4\mu \cdot 2^h)$, as $U(r, 1) - U(r - 4\alpha, 1) \leqslant 4\alpha\mu = \overline{L}/2^h$, this implies a loss of at most $\overline{L}/2^h$.

Lemma 2.26. *Let m be a man and let w be the aligned woman. Suppose we run the DA algorithm cutting at m and $r_m - \alpha$. Then the probability that every Step 2 proposal to m gives him a loss of more than L^m_α is at most* $\exp(-\alpha\beta\gamma n/2)$.

Proof. Let m be a man in $M[r_m, r_m + \delta]$. As $\mathcal{B}_3$ does not occur, m receives at least $\frac{1}{2}\alpha\beta n$ Step 2 proposals. As shown in Observation 4 of the proof sketch, each proposal gives a loss of more than L_α^m with probability at most γ. Thus, the probability that every one of these proposals give m a loss of more than L_α^m is at most

$$(1-\gamma)^{\alpha\beta n/2} \leqslant \exp(-\alpha\beta\gamma n/2).$$

□

Corollary 2. *Suppose we run the DA algorithm cutting at m and $r_m - \alpha$. Let $\mathcal{B}_6^h$ be the event that at least $2[1+(n-1)\delta]\cdot\exp(-\alpha\beta\gamma n/2)$ men in $M[r_m, r_m+\delta]$ suffer loss greater than L_α^m, where $0<\delta\leqslant 1$ and $\alpha = \overline{\sigma}/(4\cdot 2^h)$. If $r_w \geqslant \overline{\sigma}$, $\beta=\gamma=\alpha\rho$, and none of $\mathcal{B}_1$–$\mathcal{B}_3$ occur, then $\mathcal{B}_6^h$ occurs with probability at most $n^{-(c+2)}$, where $\frac{1}{3}\log\left(\frac{(c+2)\ln n}{\ln[\delta n/(3(c+2)\ln n)]}\right) \leqslant h \leqslant \frac{1}{3}\log[(c+2)\ln n]$.*

Proof. Consider a man m in $M[r_m, r_m+\delta]$. By Lemma 2.26 it follows that the probability that m experiences a loss of more than L_α^m is at most $\exp(-\alpha\beta\gamma n/2)$.

This bound depends only on m's private scores for the Step 2 proposals made to him. Thus the outcomes for the different men in $M[r_m, r_m+\delta]$ are independent.

In expectation, at most $(1+(n-1)\delta)\cdot\exp(-\alpha\beta\gamma n/2)$ men in $M[r_m, r_m+\delta]$ suffer a loss of more than L_α^m, and by a Chernoff bound, at most $2(1+(n-1)\delta)\cdot\exp(-\alpha\beta\gamma n/2)$ men suffer such a loss with probability $\exp(-\delta n\cdot\exp(-\alpha\beta\gamma n/2)/3)$.

Now $\exp(-\alpha\beta\gamma n/2) = \exp\left(-\overline{L}^3 n\rho^2/[128\mu^3 2^{3h}]\right) = \exp\left(-(c+2)\ln n/2^{3h}\right)$. Let $h = \frac{1}{3}\log[(c+2)\ln n] - \frac{1}{3}\log g$. Then $\exp(-\alpha\beta\gamma n/2) = \exp(-g)$. In sum, at most $2(1+(n-1)\delta)/\exp(g)$ men in $M[r_m, r_m+\delta]$ have a loss of more than L_α^m with probability at least $1-\exp(-\delta n\cdot\exp(-g)/3)$. So the failure probability is at most $n^{-(c+2)}$ if $g \leqslant \ln(\delta n/[3(c+2)\ln n])$. □

Proof. (Of Theorem 2.16) We apply Corollary 2 with $\delta=1$. Over all the men and women whose aligned partners have public score at least $\overline{\sigma}$, this yields that at most $2n\cdot\exp(-\alpha\beta\gamma n/2) = 2n\cdot\exp\left(-(c+2)\ln n/2^{3h}\right)$ men suffer a loss of more than $\overline{L}/2^h$, and likewise for the women, with failure probability at most $2n^{-(c+1)}\log\log n$, for integer h in the range stated in the lemma.

Applying the prior analysis, if none of $\mathcal{B}_1$–$\mathcal{B}_3$ occur, then the outcome is a stable matching with the bounds on the losses as stated in the previous paragraph. For large enough n, the failure probability will be at most n^{-c}. □

2.10 Lower Bounds

2.10.1 A Lower Bound in the Linear Model

The following theorem shows that the upper bound we obtained is the best possible up to a constant factor. The intuition is as follows: the expected number of acceptable edges per agent is $\Theta(\ln n)$, excluding the agents with public ratings of less than L. So long as the constant is small enough, the variance in the number of these edges over all the agents will be sufficient to ensure a good probability that at least one agent will have no incident acceptable edge.

For the lower bound we set $\lambda = \frac{1}{2}$. We begin by identifying and bounding the probability of some bad events, denoted by $\mathcal{B}_4$ and $\mathcal{B}_5$. We then perform an analysis for the case that $\mathcal{B}_4$ and $\mathcal{B}_5$ do not occur.

To do this, we need some additional notation. In the following lemmas, we let (m, w) and (m', w') be two pairs of men and women with equal public ranks, and suppose their public ratings are $r_m, r_w, r_{m'}, r_{w'}$, respectively. We let $x = r_m - r_{m'}$ and $y = r_w - r_{w'}$. Note that $\text{sign}(x) = \text{sign}(y)$.

Event $\mathcal{B}_4$. Let $\mathcal{E}_4$ be the following event: If $|x| \leqslant 4L$, then the number of men with public ratings in the range $[r_m, r_{m'}]$ lies in the range $(2 + x \cdot (n-2) - L(n-2), 2 + x \cdot (n-2) + L(n-2))$, and similarly for the women. $\mathcal{B}_4$ is the (bad) event that $\mathcal{E}_4$ does not occur.

Lemma 2.27. *$\mathcal{B}_4$ occurs with probability at most $2n^2 \cdot n^{-L(n-2)/12 \ln n}$.*

Proof. Not counting m and m', the expected number of men with public ratings in the range $[r_m, r_{m'}]$ is $|x|(n-2)$. By a Chernoff bound, this number lies outside the range $(|x|(n-2) - L(n-$

2), $|x|(n-2)+L(n-2))$ with probability at most $e^{-L^2(n-2)/2|x|}+e^{-L^2(n-2)/3|x|} \leqslant 2e^{-L(n-2)/12}$, as $|x| \leqslant 4L$ by assumption.

The same bound applies to the women. Now we apply a union bound to all $\frac{1}{2}n(n-1)$ pairs (m, w) and (m', w') to obtain the result. □

Event $\mathcal{B}_5$. This is the event that for some pairs (m, w) and (m', w'), either (i) $|y| = |r_{w'} - r_w| > 4L$ and the number of women in the range $[r_w, r_{w'}]$ is at most $2+3L(n-2)$, or (ii) $|x| = |r_{m'} - r_m| > 4L$ and the number of men in the range $[r_m, r_{m'}]$ is at most $2+3L(n-2)$.

Lemma 2.28. *$\mathcal{B}_5$ occurs with probability at most $n^2 \cdot n^{-L(n-2)/8\ln n}$.*

Proof. We obtain a bound in case (i). Excluding w and w', the expected number of women in the range $[r_w, r_{w'}]$ is at least $|y|(n-2)$. By a Chernoff bound it is at most $y(n-2) - (|y| - 3L) \cdot (n-2)$ with probability at most $\exp\bigl(-(|y|-3L)^2(n-2)/2|y|\bigr) \leqslant \exp\bigl(-(|y|/4)^2(n-2)/2|y|\bigr) \leqslant \exp(-L(n-2)/8)$.

The same bound holds in case (ii). Now we apply a union bound to all $\frac{1}{2}n(n-1)$ pairs (m, w) and (m', w') to obtain the result. □

Theorem 2.29. *If $n \geqslant 32,000$ and $L = \frac{1}{8}(\ln n/n)^{1/3}$ then with probability at least $\frac{1}{4}n^{-1/8}$ there is no perfect matching, let alone stable matching, in which every edge is $(L, \frac{3}{2}L)$-acceptable.*

Proof. Suppose that $\mathcal{B}_4$ and $\mathcal{B}_5$ do not occur. Then, we will show that the expected number of women with no acceptable incident edge is greater than or equal to $\frac{1}{2}n^{7/8}$. As there are n women, it immediately follows that with probability at least $\frac{1}{2}n^{-1/8}$ there is no matching in which every edge is acceptable. The result now follows if the probability of $\mathcal{B}_4 \cup \mathcal{B}_5$ is at most $\frac{1}{4}n^{-1/8}$, i.e. that $2n^2 \cdot n^{-L(n-2)/12\ln n} + n^2 \cdot n^{-L(n-2)/8\ln n} \leqslant \frac{1}{4}n^{-1/8}$; $n \geqslant 32,000$ suffices.

Lemma 2.34 below shows that in expectation there are at least $n^{7/8}$ women (and men) such that every possible proposal to one of these women would cause at least one of the two parties a loss greater than L. Recall that every edge to a woman with a public rating less than $\frac{3}{2}L$ is woman-acceptable. Let w' be the topmost such woman (i.e. the one with the highest public rating). Let

w be the woman with the lowest public rating equal to or greater than $r_{w'} + 2L$, and let m be w's aligned partner. Then the edge (m, w') gives m a loss greater than L, and thus every edge (m, w'') that is man-acceptable to m will be acceptable to w'' only if it gives w'' a loss of at most L. So for men with public rating at least r_m, an edge is acceptable only if it gives both partners a loss of at most L. We show that there are at most $5Ln$ such men if $\mathcal{B}_4$ does not occur. For if this event does not occur, then the number of women in the range $[r_{w_n}, r_w]$, and hence the number of men in the range $[r_{m_n}, r_m]$, is at most $2 + \frac{9}{2}L(n-2) \leqslant 5Ln$, if $2 \leqslant \frac{1}{16}(n-2)\ln^{1/3} n/n^{1/3}$; $n \geqslant 256$ suffices. The same bound applies to the women.

Thus, there are at least $n^{7/8} - 5Ln$ women who do not have acceptable matches. So long as $\frac{1}{2}n^{7/8} \geqslant 5Ln = \frac{5}{16}n^{2/3}\ln^{1/3} n$, this implies that the number of women with no acceptable match is at least $\frac{1}{2}n^{7/8}$. This condition holds when $n \geqslant 1$. □

Lemma 2.30. *Suppose that $\mathcal{B}_4$ and $\mathcal{B}_5$ do not occur. Further suppose that either $|x| \leqslant 2L$ or $|y| \leqslant 2L$. Then, $\left|x - y\right| < 2L$.*

Proof. We consider the case that $x \leqslant 2L$. The proof for the other case is symmetric.

Let h be the the number of men in the range $[r_m, r_{m'}]$; h is also the number of women in the range $[r_w, r_{w'}]$. By Lemma 2.27, $h \in 2 + \big(|x|(n-2) - L(n-2), |x|(n-2) + L(n-2)\big)$ and if $|y| \leqslant 4L$, $h \in 2 + \big(|y|(n-2) - L(n-2), |y|(n-2) + L(n-2)\big)$. Consequently $|x - y| < 2L$ if $|y| \leqslant 4L$.

If $|y| > 4L$, as $\mathcal{B}_5$ does not occur, the number of women in the range $[r_w, r_{w'}]$ is more than $2+3L(n-2)$. But as $\mathcal{B}_4$ does not occur, the number of men is at most $2+3L(n-2)$. These numbers are supposed to be equal, and therefore $|y| > 4L$ cannot happen. □

Lemma 2.31. *Suppose that $\mathcal{B}_4$ and $\mathcal{B}_5$ do not occur. Further suppose that $y = r_w - r_{w'} \geqslant 0$. Then, the probability that edge (m, w') causes a loss of at most L to both m and w' is at most $(2L-y)\cdot(4L+y) \leqslant 8L^2$. A symmetric bound applies if $\overline{x} = r_{m'} - r_m \geqslant 0$.*

Proof. We show the proof for the first bound. The argument for the second bound is identical. m has a loss of at least y on edge (m, w'). Therefore, for (m, w') to be acceptable to m, we need

$y \leqslant 2L$. The probability that (m, w') is acceptable is $(2L - y) \cdot (2L + x)$, and by Lemma 2.30, this is at most $(2L - y) \cdot (2L + y + 2L)$. □

Lemma 2.32. *Consider an edge (m, w'). If $x < -2L$ or $y > 2L$ then (m, w') is not acceptable.*

Proof. If $y > 2L$ then m has a loss of more than $2L$, and if $x < -2L$ then w' has a loss of more than $2L$. □

Definition 13. *Let (m, w') be an edge. If $x \geqslant -2L$ and $y \leqslant 2L$ we say (m, w') passes the public rating test, and otherwise it fails the test.*

Lemma 2.33. *Suppose that $\mathcal{B}_4$ and $\mathcal{B}_5$ do not occur. Then apart from at most $3 + 6L(n-2)$ edges (m, w') all other edges incident on m fail the public rating test.*

Proof. Let w' be the lowest rated woman in $W[r_w - 2L, r_w]$. By Lemma 2.27, there are at most $2 + 3L(n-2)$ women in $W[r_{w'}, r_w] \subseteq W[r_w - 2L, r_w]$. By Lemma 2.32, for any woman w'' with a lower rating than w', (m, w'') will fail the public rating test (as for w'', $y > 2L$).

Now let m' be the highest rated man in $M[r_m, r_m + 2L]$. By Lemma 2.27, for any woman w'' with a higher rating than w', (m, w'') will fail the public rating test (as for w'', $x < -2L$). By Lemma 2.27, there are at most $2 + 3L(n-2)$ men in $M[r_m, r_m + 2L]$, and therefore there are at most $2 + 3L(n-2)$ women in $W[r_w, r_{w'}]$.

w belongs to both these sets of women. So there are at most $3 + 6L(n-2)$ women who pass the public rating test. □

Lemma 2.34. *Suppose that $\mathcal{B}_4$ and $\mathcal{B}_5$ do not occur and $n \geqslant 100$. If $L \leqslant \frac{1}{8}(\ln n/n)^{1/3}$, and there are equal numbers of men and women, then the expected number of unmatched men (and women) is at least $n^{7/8}$.*

Proof. Consider an arbitrary man m. By Lemma 2.31, each edge which passes the public rating test is acceptable with probability at most $8L^2$. By Lemma 2.33, there are at most $3 + 6L(n-2)$

such edges incident on m. Therefore the probability that all n edges incident on m cause one or both parties a loss of more than L is at least

$$\left(1-8L^2\right)^{3+6L(n-2)} = e^{(3+(6L(n-2))\ln\left(1-8L^2\right)} \geqslant n^{-1/8},$$

as we argue next.

For this to hold, it suffices that

$$
\begin{aligned}
& -8\cdot\frac{3+6L(n-2)}{\ln n}\left(-8L^2-\frac{1}{2}(8L^2)^2-\frac{1}{3}(8L^2)^3-\ldots\right)\leqslant 1 \\
\text{or that}\quad & 8\cdot\frac{7Ln}{\ln n}\cdot\frac{8L^2}{1-8L^2}\leqslant 1 \qquad (\text{if } 3\leqslant Ln) \\
\text{or that}\quad & 56\cdot\frac{n^{2/3}}{8\ln^{2/3}n}\cdot\frac{\ln^{2/3}n}{8n^{2/3}-\ln^{2/3}n}\leqslant 1 \\
\text{or that}\quad & \frac{7}{8}\cdot\frac{1}{1-\ln^{2/3}n/8n^{2/3}}\leqslant 1
\end{aligned}
$$

which holds if $\ln^{2/3} n/8n^{2/3} \leqslant 1/8$, which holds for $n \geqslant 1$. Our other condition, $3 \leqslant Ln$, or $3 \leqslant \frac{1}{8}n^{2/3}\ln^{1/3} n$, holds if $n \geqslant 100$.

Thus the expected number of women having all incident edges causing a loss of more than L to both parties is at least $n^{7/8}$. □

2.10.2 Lower Bound on Performance for the General Utility Model

Now we show that without the bounds on the derivatives, no sub-constant loss is achievable in general.

Definition 14 (Sub-constant function)**.** *A function $f(x) : \mathbb{R} \rightarrow \mathbb{R}^+$ is sub-constant if for every choice of constant $c > 0$, there exists an $\overline{x}$ such that for all $x \geqslant \overline{x}$, $f(x) \leqslant c$.*

We first examine what happens if the derivatives w.r.t. private scores are not bounded, but the derivatives w.r.t. public ratings are bounded; this implies there is no lower bound on the ratio of the derivatives of the utility functions w.r.t. public ratings and private scores (recall Definition 5).

Lemma 2.35. *Let $f : N \longrightarrow \mathbb{R}^+$ be a continuous, strictly decreasing sub-constant function, and let $\sigma, \delta \in (0, 1)$ be constants. Suppose the public ratings and private scores of the n men and n women are drawn uniformly and independently from $[0, 1]$. Then there exist continuous and strictly increasing utility functions $U(.,.)$ and $V(.,.)$ having derivatives w.r.t. their first variables that are bounded by a constant, but for (at least) one of which the derivatives w.r.t. their second variables are not bounded by any constant, having the following property: for some $\overline{n} > 0$, for all $n \geqslant \overline{n}$, with probability at least $1 - \delta$, in every perfect matching, some rank i man m_i or woman w_i with public rating at least σ receives utility less than $U(r_{w_i}, 1) - f(n)$ or $V(r_{m_i}, 1) - f(n)$, respectively.*

Proof. We will give an example where, with probability at least $1 - \delta$, in every perfect matching, man m_1 receives utility less than $U(r_{w_1}, 1) - f(n)$.

Observe that proving the result for a more slowly decreasing f implies it for faster decreasing functions. In what follows, at times we will need to assume f decreases sufficiently slowly, but given the just made observation, we can do so WLOG.

Now we define $U(r, s) = r + g(s)$ where $g(s)$ is a continuous, strictly increasing function, and for which

$$g\Big(1 - \frac{\delta}{8n \cdot f(n)}\Big) = g(1) - f(n).$$

The reason for this condition will become clear in due course. We will first demonstrate that there is such a g. To this end, define

$$\begin{aligned} k(y) &= \frac{8y \cdot f(y)}{\delta} && \text{for } y \geqslant 0 \\ g(s) &= g(1) - f(k^{-1}(1/(1-s))), && \text{if } s < 1 \\ g(1) &= f(k^{-1}((1))) && \text{(so } g(0) = 0). \end{aligned}$$

We will want k to be strictly increasing and unbounded. This is true if f is sufficiently slowly decreasing. Next, note that as f is continuous, so is k. Therefore k^{-1} is continuous and strictly increasing, and therefore so is g, except possibly at $s = 1$. For g to be continuous at $s = 1$ we need $\lim_{s\to 1} f(k^{-1}(1/(1-s))) = 0$, which happens as $\lim_{x\to\infty} f(x) = 0$, which happens since f is sub-constant.

Setting $s = 1 - \frac{\delta}{8n\cdot f(n)}$ gives

$$g(s) = g(1) - f\left(k^{-1}\left(\frac{8n \cdot f(n)}{\delta}\right)\right) = g(1) - f(n), \tag{2.4}$$

as desired.

Strictly speaking, we should rescale the utility so that its range is $[0,1]$ rather than the actual $\left[0, 1 + g(1)\right] \subset \left[0, 1 + f(0)\right]$. Note that although $1 + g(1)$ is a function of δ, it is always bounded by $1 + f(0)$, a constant, and so the rescaling does not affect the result stated in the lemma. We omit performing the rescaling to avoid unnecessary clutter.

For m_1 to face a loss of at most $f(n)$, he must match with a woman having public rating at least $r_{w_1} - f(n)$.

The probability that no woman has a public rating in the range $[1 - \ln(4/\delta)/n, 1]$ is at most

$$\left(1 - \frac{\ln(4/\delta)}{n}\right)^n \leqslant \exp(-\ln(4/\delta)) = \tfrac{1}{4}\delta.$$

Otherwise, $r_{w_1} \geqslant 1 - \ln(4/\delta)/n$. If f is sufficiently slowly decreasing, for large enough n, $\ln(4/\delta)/n \leqslant f(n)$. Therefore, for such large enough n, with probability at least $1 - \frac{1}{4}\delta$, $r_{w_1} - f(n) \geqslant 1 - 2f(n)$. Call the probability $\frac{1}{4}\delta$ event $\mathcal{B}_1$.

The same analysis shows that, with failure probability at most $\frac{1}{4}\delta$, $r_{m_1} \geqslant 1 - f(n)$, and as f is a sub-constant function, for large enough n, $r_{m_1} \geqslant 1 - f(n) \geqslant \sigma$. Call the probability $\frac{1}{4}\delta$ event $\mathcal{B}_2$.

The expected number of women other than w_1 in $W[r_{w_1} - f(n), r_{w_1}]$ is $(n-1)\cdot f(n)$. Let n_w

be the actual number of women other than w_1 in this range. By a Chernoff bound,

$$\Pr\left[n_w \geqslant (n-1)\cdot f(n) + \sqrt{3(n-1)\cdot f(n)\cdot \ln(4/\delta)}\right] \leqslant e^{-\ln(4/\delta)} = \tfrac{1}{4}\delta.$$

If f is decreasing sufficiently slowly, then for sufficiently large n, $1+\sqrt{3(n-1)\cdot f(n)\cdot \ln(4/\delta)} \leqslant n\cdot f(n)$; therefore, in addition, $1+n_w \leqslant 2n\cdot f(n)$ with probability at least $1-\frac{1}{4}\delta$. Note that $1+n_w$ is the number of women in $W[r_{w_1} - f(n), 1]$. Call the probability $\frac{1}{4}\delta$ event $\mathcal{B}_2$.

Next, consider an edge (m_1, w) for which m_1's private score is s. For this edge to cause more than $f(n)$ loss to m, it suffices that $g(s) < g(1) - f(n) = g(1-\delta/[8n\cdot f(n)])$ by (2.4). This occurs with probability $\delta/[8n\cdot f(n)]$.

We now lower bound the probability that every match with a woman in $W[r_{w_1} - f(n), 1]$ causes m a loss of more than $f(n)$ if none of the events $\mathcal{B}_1$–$\mathcal{B}_3$ occur. This probability is at least:

$$\left(1-\frac{\delta}{8n\cdot f(n)}\right)^{2n\cdot f(n)} \geqslant 1-\tfrac{1}{4}\delta.$$

Thus, by a union bound, modulo an overall failure probability of at most δ, m_1 has a loss of more than $f(n)$ on every incident edge, and hence in every perfect matching some agent (m_1 actually) incurs a loss of more than $f(n)$. □

We now consider the case where the derivatives w.r.t. the first variable are bounded, but there is no bound on the derivatives w.r.t. the second variable.

Lemma 2.36. *Let $f : N \to \mathbb{R}^+$ be a continuous, strictly decreasing sub-constant function, and let $\sigma, \delta \in (0,1)$ be constants. Suppose the public ratings and private scores of the n men and n women are drawn uniformly and independently from $[0,1]$. Then there exist continuous and strictly increasing utility functions $U(.,.)$ and $V(.,.)$ having derivatives w.r.t. their second variables that are bounded by a constant, but for (at least) one of which the derivative w.r.t. their first variable is not bounded by any constant, having the following property: for some $\overline{n} > 0$, for all $n \geqslant \overline{n}$, with probability at least*

$1-\delta$*, in every perfect matching, some man* m_i *or woman* w_i *with public rating at least* σ *receives utility less than* $U(r_{w_i}, 1) - f(n)$ *or* $V(r_{m_i}, 1) - f(n)$*, respectively.*

Proof. We will give an example where, in every stable matching, man m_1 receives utility less than $U(r_{w_1}, 1) - f(n)$. The analysis has the same thrust as the one for the preceding lemma.

Let $U(r, s) = \widetilde{g}(r) + s$, where $\widetilde{g}$ is defined below is a very similar way to the g in the proof of Lemma 2.35.

$$\begin{aligned}
\widetilde{k}(y) &= \frac{8y \cdot f(y)}{\delta} \\
\widetilde{g}(r) &= \widetilde{g}(1) - 3 \cdot f(\widetilde{k}^{-1}(1/(1-r))), && \text{if } r < 1 \\
\widetilde{g}(1) &= 3 \cdot f(\widetilde{k}^{-1}((1))) && (\text{so } \widetilde{g}(0) = 0).
\end{aligned}$$

Now, setting $r = 1 - \frac{\delta}{8n \cdot f(n)} \triangleq 1 - v$, gives $\widetilde{g}(r) = \widetilde{g}(1) - 3 \cdot f(n)$. Again, strictly speaking, we should rescale the utility so that its range is $[0, 1]$.

Next, we will show that $\widetilde{g}(1 - v/3) - \widetilde{g}(r) \geqslant f(n)$. As $\widetilde{g}(r) = \widetilde{g}(1) - 3 \cdot f(n)$, it suffices to show that $\widetilde{g}(1 - v/3) \geqslant \widetilde{g}(1) - 2 \cdot f(n)$, which we do as follows:

$$\begin{aligned}
\widetilde{g}(1 - v/3) &= \widetilde{g}(1) - 3 \cdot f\Big(\widetilde{k}^{-1}\Big(\frac{3}{v}\Big)\Big) \\
&= \widetilde{g}(1) - 3 \cdot f\Big(\widetilde{k}^{-1}\Big(3 \cdot \frac{8n \cdot f(n)}{\delta}\Big)\Big) \\
&= \widetilde{g}(1) - 3 \cdot f\Big(\widetilde{k}^{-1}\Big(\frac{3}{2} \cdot \frac{f(n)}{f(2n)} \cdot \frac{8 \cdot 2n \cdot f(2n)}{\delta}\Big)\Big)
\end{aligned}$$

If f is sufficiently slowly decreasing, for all n, $\frac{3}{2} \cdot \frac{f(n)}{f(2n)} \geqslant 1$, and as $\widetilde{k}^{-1}$ is increasing and f is decreasing, the RHS of the above expression is at least

$$\widetilde{g}(1) - 3 \cdot f\Big(\widetilde{k}^{-1}\Big(\frac{8 \cdot 2n \cdot f(2n)}{\delta}\Big)\Big) = \widetilde{g}(1) - 3 \cdot f(2n) \geqslant \widetilde{g}(1) - 2 \cdot f(n),$$

as $3 \cdot f(2n) \geqslant 2 \cdot f(n)$, if f decreases sufficiently slowly.

The probability that no woman has a public rating in the range $[1 - \ln(4/\delta)/n, 1]$ is at most $\frac{1}{4}\delta$. As f is a sub-constant function, for large enough n, $\ln(4/\delta)/n \leqslant \frac{1}{3} \cdot \delta/[8n \cdot f(n)] = \frac{1}{3}\nu$.

Therefore, for large enough n, with probability at least $1 - \frac{1}{4}\delta$,

$$\widetilde{g}(r_{w_1}) - \widetilde{g}(1-\nu) \geqslant \widetilde{g}(1-\nu/3) - \widetilde{g}(1-\nu) \geqslant f(n).$$

Thus, with probability at least $1 - \frac{1}{4}\delta$, all women with public rating less than $(1-\nu)$ will cause m_1 a loss of more than $f(n)$. Call the probability $\frac{1}{4}\delta$ event $\mathcal{B}_1$.

Let n_w be the actual number of women, aside w_1, with public rating at least $(1-\nu)$. $\mathrm{E}\left[n_w\right] \leqslant (n-1)\nu$. By a Chernoff bound,

$$\Pr\left[n_w \geqslant (n-1)\nu + \sqrt{3(n-1)\nu \cdot \ln(4/\delta)}\right] \leqslant \exp(-\ln(4/\delta)) = \tfrac{1}{4}\delta.$$

However slowly f is decreasing, for large enough n, $1 + \sqrt{3(n-1)\nu \cdot \ln(4/\delta)} \leqslant n\nu$, which implies $1 + n_w \leqslant 2n\nu$. Call the probability $\frac{1}{4}\delta$ event $\mathcal{B}_2$.

The same analysis as in the proof of Lemma 2.35 shows that, with failure probability at most $\frac{1}{4}\delta$, $r_{m_1} \geqslant 1 - f(n)$, and as f is a sub-constant function, for large enough n, $r_{m_1} \geqslant 1 - f(n) \geqslant \sigma$. Call the probability $\frac{1}{4}\delta$ event $\mathcal{B}_3$.

Next, note that an edge (m_1, w) causes m_1 a loss of more than $f(n)$ based on the private score alone with probability $1 - f(n)$.

Thus, if none of the events $\mathcal{B}_1$–$\mathcal{B}_3$ occur, the probability that every edge incident on m_1 causes it a loss of more than $f(n)$ is at least

$$\left(1 - f(n)\right)^{2n\nu} = \left(1 - f(n)\right)^{\delta/[4 \cdot f(n)]} \geqslant 1 - \delta/4,$$

if $\delta \leqslant 1$.

Therefore, by a union bound, modulo an overall failure probability of at most δ, m_1 has a loss of more than $f(n)$ on every incident edge, and hence in every perfect matching some agent (m_1 actually) incurs a loss of more than $f(n)$. □

2.11 Epsilon-Bayes-Nash Equilibria

In this section, we demonstrate that there is a ε-Nash equilibrium in which with high probability all agents have low losses. To obtain this result, we need the stronger bounded-derivatives condition, namely we need both lower and upper bounds for the two derivative expressions (see Definition 8). We will assume that both U and V satisfy the strong bounded derivative property.

Our analysis here will repeatedly use weak stochastic dominance to justify the application of Chernoff bounds. To avoid repetition, we summarize the technique here.

We begin by proving a useful technical lemma.

Lemma 2.37. *Suppose* $\mathcal{X} = \{X_1, X_2, \ldots, X_m\}$ *is a collection of not-necessarily independent binary random variables. Suppose that* $\Pr[X_i = 1 | \mathcal{X} \setminus \{X_i\}] \leqslant p_i$. *Let* Y_i *be a binary variable with* $\Pr[Y_i = 1] = p_i$, *with the* Y_i *being independent. For all* z, $\Pr[\sum_i X_i \geqslant z] \leqslant \Pr[\sum_i Y_i \geqslant z]$.

Proof. We will define a collection of random variables $\mathcal{Z} = \{Z_1, Z_2, \ldots, Z_m\}$ such that Z_i has the exact same distribution as Y_i and whenever $\sum_i X_i \geqslant z$, then $\sum_i Z_i \geqslant z$.

Draw the X_i sequentially. Suppose $X_1 = x_1, X_2 = x_2, \ldots X_{i-1} = x_{i-1}$.
Let $\Pr[X_i = 1 | X_1 = x_1, X_2 = x_2, \ldots, X_{i-1} = x_{i-1}] = q_i \leq p_i$. We define Z_i as follows:

- If $X_i = 1$, set $Z_i = 1$.
- If $X_i = 0$, set $Z_i = 1$ with probability $\frac{p_i - q_i}{1 - q_i}$ and set $Z_i = 0$ with probability $\frac{1 - (p_i - q_i)}{1 - q_i}$.

Notice that $\Pr[\sum_i X_i \geqslant z] \leqslant \Pr[\sum_i Z_i \geqslant z] = \Pr[\sum_i Y_i \geqslant z]$, which completes the proof.

□

By Lemma 2.37, for all z, $\Pr[\sum_i X_i \geqslant z] \leqslant \Pr[\sum_i Y_i \geqslant z]$. Thus, if by means of a Chernoff bound, we show that $\Pr[\sum_i Y_i \geqslant z] \leqslant q$, then $\Pr[\sum_i X_i \geqslant z] \leqslant q$ also. Henceforth, we will justify this application of a Chernoff bound to $\mathcal{X}$ by saying it uses *stochastic dominance.*

Our analysis will build on the bound shown in Theorem 2.20. It will be helpful to review the randomness that was used. The probability that events $\mathcal{B}_1$ and $\mathcal{B}_2$ do not occur is based on the public ratings of the men and women. The bound on the probability that event $\mathcal{B}_3$ occurs for a particular man m_i is based on the private scores of the edges to m_i, namely the private score of the woman w_j for the man m_i, for each such edge (m_i, w_j). Note that $\mathcal{B}_3$ is the bad event in Claim 2.4.4. The bound on the final error term is based on the private scores of m_i for their edges to the women w_j, as given in Claim 2.4.5. We will let $\mathcal{B}_7$ denote the bad event in Claim 2.4.5, namely that all the proposals to m_i cause him too large a loss. Symmetric bounds apply to the women.

In the analysis that follows, we will identify additional bad events concerning there being too few or too many agents in a range of public ratings; these will depend on the range. We will also bound the probability of losses for bottommost women and men using the private scores of proposals to these agents; these private scores will be disjoint from the ones used in the bounds mentioned in the previous paragraph.

In the remainder of this section, m and w are always aligned, as are m' and w', m'' and w'', etc.

In addition, in order to improve some of the bounds, we will restate losses in terms of public ratings and private scores. A quick inspection of the proof of Theorem 2.15 shows that the high probability bound on the loss for a man m, whose aligned woman w has public rating $r_w \geqslant \overline{\sigma}$, is at most $U(r_w, 1) - (r_w - \overline{\sigma}, 1)$ (recall that $\alpha = \frac{1}{4}\overline{\sigma}$). Also note that is suffices to set $\overline{\sigma} = [128(c+2)\ln n/(\rho_l^2 n)]^{1/3}$. Similarly, if $r_w \geqslant \overline{\sigma}/t$, where $t > 1$, the bound on the loss is at most $L_t^m \triangleq U(r_w, 1) - U(r_w - \overline{\sigma}t^2, 1)$ (see the proof of Theorem 2.17). Analogously, for a woman w, if $r_m \geqslant \overline{\sigma}/t$, where $t > 1$, the bound on the loss is at most $L_t^w \triangleq V(r_m, 1) - V(r_m - \overline{\sigma}t^2, 1)$.

As already noted, these bounds use the private scores of proposals to men and women with public ratings of at least $\overline{\sigma}/t$.

Shortly, we will specify maximum values t_m and t_w of t for the men and women, respectively. We will demonstrate the existence of a stable match in which w.h.p. every man m has a loss of at most $L^m_{t_m}$, and every woman w has a loss of at most $L^w_{t_w}$. For this to be meaningful when $r_w - \overline{\sigma}t_m^2 < 0$, we extend the definition of U to this domain as follows. For $r < 0$, $\frac{\partial U(r,s)}{\partial r} = \mu_\ell$ and $\frac{\partial U(r,s)}{\partial s} = \rho_\ell$. We proceed analogously to handle the case that $r_m - \overline{\sigma}t_w^2 < 0$.

We define t_m and t_w using suitable constants $\eta > 1$ and $0 < \nu < 1$, which we will specify later. We set $\sigma_m = \nu/n^{1/3}$ and $\sigma_w = \eta/n^{1/3}$. We then define $t_m = \overline{\sigma}/\sigma_m$ and $t_w = \overline{\sigma}/\sigma_w$. Note that $\sigma_w/\sigma_m = \eta/\nu$.

The maximum loss will occur only to some of the agents with low public ratings. We identify the potentially high-loss agents as follows.

Definition 15. *Let w' be the bottommost woman with a public rating of at least σ_m and let m' be aligned with w'. Then the* bottom zone *of men comprises the set $B_M \triangleq M[0, r_{m'})$, and the top zone T_M comprises $M[r_{m'}, 1]$. Similarly, let m'' be the bottommost man with a rating of at least σ_w and let w'' be aligned with m''. Then the* bottom zone *of women comprises the set $B_W \triangleq W[0, r_{w''})$, and the top zone T_W comprises $W[r_{w''}, 1]$.*

Note that by Theorem 2.17, in any stable match, every man $m \in T_M$ has loss at most $L^m_{t_m}$. Likewise, every woman $w \in T_W$ has loss at most $L^w_{t_w}$.

We also want to distinguish those edges which yield men a utility of at least $U(0,1)$ and women a utility of at least $V(0,1)$.

Definition 16. *An edge (m_i, w_j) is* man-high *if $U(r_{w_j}, s_{m_i}(w_j)) \geqslant U(0,1)$, and otherwise it is* man-low*; it is* woman-high *if $V(r_{m_i}, s_{w_j}(m_i)) \geqslant V(0,1)$, and otherwise it is* woman-low.

We begin by identifying two bad events $\mathcal{B}_8$ and $\mathcal{B}_9$ and bounding the probabilities they occur.

Event $\mathcal{B}_8$. Let $\mathcal{E}_8$ be the event that the number of men in B_M lies in the range $(\frac{1}{2}\sigma_m \cdot n, 2\sigma_m \cdot n)$, and each of these men has public rating less than $3\sigma_m$, together with the corresponding event for women. Let $\mathcal{B}_8$ be the complementary event.

Lemma 2.38. *$\mathcal{B}_8$ occurs with probability at most* $\exp(-\sigma_m \cdot n/3) + \exp(-\sigma_m \cdot n/6) + \exp(-\sigma_m \cdot n/8) + \exp(-\sigma_w \cdot n/3) + \exp(-\sigma_w \cdot n/6) + \exp(-\sigma_w \cdot n/8) \leqslant 6\exp(-\sigma_m \cdot n/8)$. *This bound is based on the independent random choices of public ratings for the men and women.*

Proof. In expectation, there are $\sigma_m \cdot n$ women with public rating less than σ_m. These choices are based on the women's independent public scores. Hence, by a Chernoff bound, the probability that there are at least $2\sigma_m \cdot n$ women with public rating less than σ_m is at most $\exp(-\sigma_m \cdot n/3)$, and the probability that there are at most $\frac{1}{2}\sigma_m \cdot n$ women with public rating less that σ_m is at most $\exp(-\sigma_m \cdot n/8)$. But these are the women aligned with the men in B_M. Hence these bounds also apply to the number of men in B_M.

Now we bound the probability that there are at most $2\sigma_m \cdot n$ men in the public rating range $[0, 3\sigma_m)$. The expected number of men in this range is $3\sigma_m \cdot n$. This is based on their independent public ratings. Then, by a Chernoff bound, there are at most $2\sigma_m \cdot n$ men in this range with probability at most $\exp(-\sigma_m \cdot n/6)$.

Analogous bounds apply to the women. $\square$

Event $\mathcal{B}_9$. This is the event that $r_{m'} < 4\sigma_m$.

Lemma 2.39. *If $\mathcal{B}_8$ does not occur, then $\mathcal{B}_9$ occurs with probability at most* $\exp(-\sigma_m \cdot n)$. *This bound is based on the independent random choices of public ratings for the men.*

Proof. As $\mathcal{B}_8$ does not occur, by Lemma 2.38,
$B_M \subset M[0, 3\sigma_m)$. Therefore, if $r_{m'} \geqslant 4\sigma_m$, $M[3\sigma_m, 4\sigma_m)$ is empty. But the probability that $M[3\sigma_m, 4\sigma_m)$ is empty is at most $(1-\sigma_m)^n \leqslant \exp(-\sigma_m \cdot n)$, and it follows that this is the probability that $r_{m'} \geqslant 4\sigma_m$. $\square$

The desired stable match will be found by running the woman-proposing DA when each man m, whose aligned woman w has public rating less than σ_m, applies a truncation strategy of refusing proposals that provide a loss greater than $L^m_{t_m}$. No truncation is applied by men with higher public ratings, but we already know their losses are bounded by $L^m_{t_m}$. The women apply a symmetric truncation, meaning that a woman w will only propose edges that provide a loss of at most $L^w_{t_w}$.

Our analysis considers the result of running woman-proposing DA on the truncated edge set. We begin by observing that every man in T_M is matched, and similarly every woman in T_W is matched. We then argue that every woman in B_W will be matched, from which we deduce that every man in B_M must also be matched.

As we showed in Theorem 2.17, with failure probability $O(n^{-(c+1)})$, in every stable match, every man m in T_M will have a loss of at most $L^m_{t_m}$. Furthermore, this match is achieved with the edge set cut as in Lemma 2.13. As the men in T_M do not truncate any edges, all the edges required for Lemma 2.13 remain present despite the men's truncations. Also, all the edges used by this lemma are women-high, and the women do not truncate such edges. Thus the result of Lemma 2.13 continues to apply as does Theorem 2.17.

A symmetric argument shows that with failure probability $O(n^{-(c+1)})$, in every stable match, every woman w in T_W will have a loss of at most $L^w_{t_w}$.

To analyse what happens to the women in B_W we proceed as follows.

We observe that w.h.p.:

i. The men in B_M receive at most $\frac{1}{4}|B_M|$ proposals which are both man-high and woman-high.

ii. The men in B_M receive at most $\frac{1}{4}|B_M|$ proposals which are both man-high and woman-low.

iii. We conclude that at most half the men in B_M will receive a man-high proposal.

iv. The proposals from B_W to B_M that are both man-low and woman-low behave in the same way as in the uniform random model, up to a constant factor. This will mean that it suffices that the women in B_W have $\Theta(\ln^2 n)$ man-low and women-low edges to the men in B_M (which they do), and ensures that each man in B_M receives at least one proposal.

Our analysis will also be concerned with the following subsets W_h of women, for integer $h \geqslant 0$; W_h comprises the women aligned with the men in $M[2^h\sigma_w, 2^{h+1}\sigma_w)$.

Event $\mathcal{B}_{10}$. $\mathcal{B}_{10}^h$ is the event that $|W_h| \geqslant 2^{h+1}\eta \cdot n^{2/3}$. And $\mathcal{B}_{10} = \cup_{h \geqslant 0}\mathcal{B}_{10}^h$.

Lemma 2.40. *$\mathcal{B}_{10}^h$ occurs with probability at most $\exp\left(-2^h\eta \cdot n^{2/3}/3\right)$. And $\mathcal{B}_{10}$ occurs with probability at most $2\exp\left(-\eta \cdot n^{2/3}/3\right)$, if $n^{2/3} \geqslant 3$. These bounds are based on the independent random choices of public ratings for the men.*

Proof. The expected number of men in $M[2^h\eta/n^{1/3}, 2^{h+1}\eta/n^{1/3})$ is $2^h\eta n^{2/3}$, and these choices are based on the men's public ratings. Thus, by a Chernoff bound, there are at least $2^{h+1}\eta n^{2/3}$ men in this range with probability at most $\exp\left(-2^h\eta n^{2/3}/3\right)$. This is also the bound on the number of women aligned with these men.

The second claim follows on summing the probability bound over h, using the assumption that $n^{2/3} \geqslant 3$. □

Lemma 2.41. *Suppose that none of $\mathcal{B}_1$–$\mathcal{B}_{10}$ occur. Then there are at most $\frac{1}{4}|B_M|$ matches between women in T_W and men in B_M, with failure probability at most $\exp(-|B_M|/24)$, if $\eta \geqslant 6\nu$ and $4(\eta/\nu) \cdot \exp\left(-(\eta/2)^3\rho_\ell^2/128\right) \leqslant \frac{1}{10}$.*

Proof. We will consider the sets W_h of women aligned with $M[2^h\eta/n^{1/3}, 2^{h+1}\eta/n^{1/3})$, for $h \geqslant 0$. The union of the sets forms T_W.

If a woman w_j in W_h is matched to a man m_i in B_M the difference in public scores between m_j and m_i is at least

$$r_{m_j} - r_{m_i} \geqslant (2^h\eta - 3\nu)/n^{1/3} \geqslant 2^{h-1}\eta/n^{1/3} \triangleq g_h,$$

as $\eta \geqslant 6\nu$ and $\mathcal{B}_8$ does not occur (and hence $r_{m_i} \leqslant 3\sigma_m$).

We will apply Lemma 2.26, swapping the roles of the men and women, with $\alpha = \frac{1}{4}g_h$, $\beta = \gamma = \alpha\rho_\ell$, to bound the probability p_h that w_j sustains a loss of more than $V(r_{m_j}, 1) - V(r_{m_j} - g_h, 1)$. To match with any man m_i in B_M, w_j must sustain such a loss. Therefore, with probability at least $1 - p_h$, w_j does not match with a man in B_M.

As none of $\mathcal{B}_1$–$\mathcal{B}_{10}$ occur, By Lemma 2.26, the probability that in the man-proposing DA with cuts at w_j and $r_{w_j} - \alpha$, gives her a loss of more that L_h^m is at most $\exp\Big(-2^{3(h-1)}(\eta/n^{1/3})^3\rho_\ell^2 n/128\Big)$, and this bound depends only on the private scores of the proposals between the woman and the men in T_M. But if this does not occur, L_h^m is also a bound on w_j's loss in the woman proposing DA. As $L_h^m \leqslant g_h$, this implies m_j is matched to a man in T_M.

As $\mathcal{B}_{10}$ does not occur, by Lemma 2.40, $|W_h| \leqslant 2^{h+1}\eta n^{2/3}$. Also, as $\mathcal{B}_8$ does not occur, by Lemma 2.38, $|B_M| \geqslant \frac{1}{2}\sigma_m \cdot n = \frac{1}{2}\nu n^{2/3}$. Finally, recall that $\overline{\sigma}^3 = 128(c+2)\ln n/(\rho_\ell^2 n)$. Thus the expected number of matches between women in T_W and men in B_M is at most

$$(2^{h+1}\eta n^{2/3}) \cdot \exp\Big(-2^{3(h-1)}(\eta/n^{1/3})^3\rho_\ell^2 n/128\Big) \leqslant 2^{h+2}(\eta/\nu) \cdot (\tfrac{1}{2}\nu n^{2/3}) \cdot \exp\Big(-2^{3h}(\eta/2)^3\rho_\ell^2/128\Big).$$

As $4(\eta/\nu) \cdot \exp\big(-(\eta/2)^3\rho_\ell^2/128\big) \leqslant \frac{1}{10}$, we see that $\exp\big(-(\eta/2)^3\rho_\ell^2/128\big) \leqslant \nu/40\eta$, and therefore the bound on the number of matches is at most

$$\frac{2^h \cdot (\frac{1}{2}\nu n^{2/3})}{10\left(\frac{40\eta}{\nu}\right)^{2^{3h}-1}}.$$

Summing over all $h \geqslant 0$, we obtain that the expected number of matches is at most $\frac{1}{8} \cdot (\frac{1}{2}\nu n^{2/3})$. By a Chernoff bound, the number of matches is at most $\frac{1}{4} \cdot (\frac{1}{2}\nu n^{2/3}) \leqslant \frac{1}{4}|B_M|$, with failure probability at most $\exp\Big(\nu n^{2/3}/48\Big)$.

Next, we argue that this use of a Chernoff bound is justified by stochastic dominance. The expectation is the product of two terms: a bound on $|W_h|$, which follows from the assumption that $\mathcal{B}_{10}$ does not occur, and a bound on the probability that an arbitrary woman w in W_h has a small loss and therefore cannot be proposing to any man in B_M. The upper bound on the latter probability depends only on the men's and women's private scores for the proposals from w to the men in T_M, and so we can safely apply stochastic dominance. □

Lemma 2.42. *Suppose that that neither $\mathcal{B}_8$ nor $\mathcal{B}_9$ occur. Then, the probability that a proposal from a woman in B_W to a man in B_M is man-high is at most $\eta \nu^2 \rho_u \rho_\ell / [32(c+2)\ln n]$, if $t_m \geqslant 2$.*

Proof. Since $\mathcal{B}_8$ does not occur, by Lemma 2.38, every woman in B_W has rating at most $3\sigma_w$. We now use this to bound the probability that an edge from woman $w_j \in B_W$ to man $m_i \in B_M$ is man-high. For the edge to be man-high, we need $U(r_{w_j}, s_{m_i}(w_j)) \geqslant U(0,1)$. Now, $U(0, s_{m_i}(w_j) + r_{w_j} \cdot \rho_u) \geqslant U(r_{w_j}, s_{m_i}(w_j))$, so the edge is man-high with probability at most $r_{w_j} \cdot \rho_u \leqslant 3\sigma_w \cdot \rho_u$.

Because of the truncation, the edge is low if $U(r_{w_i} - \overline{\sigma} \cdot t_m^2, 1) \leqslant U(r_{w_j}, s_{m_i}(w_j)) < U(0,1)$. Because $r_{w_i} \leqslant \sigma_m = \overline{\sigma}/t_m$, the edge is low if $U(\overline{\sigma}/t_m - \overline{\sigma} \cdot t_m^2, 1) \leqslant U(r_{w_j}, s_{m_i}(w_j)) < U(0,1)$. Consequently the probability that the edge is low is at least $\rho_\ell(\overline{\sigma} t_m^2 - \overline{\sigma}/t_m) \geqslant \frac{3}{4}\rho_\ell \overline{\sigma} t_m^2$, as $t_m \geqslant 2$.

Therefore the probability that a proposal is man-high is at most

$$\frac{3\sigma_w \cdot \rho_u}{\frac{3}{4}\overline{\sigma} t_m^2 \rho_\ell} = \frac{4\sigma_w \cdot \rho_u}{\overline{\sigma} t_m^2 \rho_\ell}.$$

Recall that $\overline{\sigma} = \sigma_m \cdot t_m$ and $\overline{\sigma}^3 = 128(c+2)\ln n/(\rho_l^2 n)$. Thus, the probability bound is

$$\frac{4\sigma_w \cdot \rho_u \cdot \sigma_m^2}{\overline{\sigma}^3 \rho_\ell} = \frac{4\eta \nu^2 \rho_u \rho_\ell}{128(c+2)\ln n} = \frac{\eta \nu^2 \rho_u \rho_\ell}{32(c+2)\ln n}.$$

□

We will now analyze the women-low proposals. Note that once a women makes one such proposal, all her subsequent proposals will be woman-low. We now state two assumptions regarding the proposals by women in B_W. They will be demonstrated later.

Assumption 1. *i. The edges proposed by each woman in B_W have private score at least $\frac{1}{2}$.*
ii. Each woman in B_W proposes to at most half the men in B_M.

Lemma 2.43. *Let w be a woman in B_W, who is now proposing woman-low edges. For her next proposal, let $p_{\min}$ be the minimum probability that she selects a particular man in B_M, and let $p_{\max}$ be the maximum probability, over the men she has not yet proposed to. Then $p_{\max}/p_{\min} \leqslant 2\mu_u/\mu_\ell$.*

Proof. Suppose w's most recent proposal provided her a utility of u. Consider the utility interval $(u, u - \delta u]$. The probability that she selects a man providing utility in this interval is given by the private score decrease that reduces the utility u to $u - \delta u$ divided by the remaining available private score, which includes the range $[0, \frac{1}{2}]$ by assumption. Thus the probability that she selects a particular man in B_M varies between $\delta u \cdot \mu_\ell$ and $2\delta u \cdot \mu_u$. □

Corollary 3. *There are at most $d \ln n|B_M|$ woman-low proposals to men in B_M, where $d = 2(c + 2)\mu_u/\mu_\ell$, with failure probability at most $n^{-(c+1)}$.*

Proof. Suppose $m \in B_M$ does not receive a woman-high proposal. The probability that m receives no proposals among $d \ln n|B_M|$ woman-low proposals is at most

$$\left(1 - \frac{\mu_\ell}{2\mu_u|B_M|}\right)^{d \ln n|B_M|} \leqslant \exp(-d\mu_\ell \ln n/2\mu_u).$$

As $d = 2(c+2)\mu_u/\mu_\ell$, the probability is at most $n^{-(c+2)}$. A union bound over the men in B_M yields the claim. □

Lemma 2.44. *The number of man-high proposals from women in B_W to men in B_M is at most $d\eta v^2\rho_u\rho_\ell|B_M|/[16(c+2)]$, with failure probability at most $\exp\big(-d\eta v^2\rho_u\rho_\ell|B_M|/[48(c+2)]\big)$.*

Proof. By Lemma 2.42, the probability that a proposal is man-high is at most $\eta v^2\rho_u\rho_\ell/[32(c + 2)\ln n]$. Over $d \ln n|B_M|$ proposals, this yields an expected $d\eta v^2\rho_u\rho_\ell|B_M|/[32(c + 2)]$ proposals. By a Chernoff bound, there are at most $d\eta v^2\rho_u\rho_\ell|B_M|/[16(c + 2)]$ such proposals with failure probability at most $\exp\big(-d\eta v^2\rho_u\rho_\ell|B_M|/[48(c+2)]\big)$. □

Lemma 2.45. *Suppose neither $\mathcal{B}_8$ nor $\mathcal{B}_9$ occur. Then, over the course of the first $d \ln n|B_M|$ woman-low proposals from women in B_W to men in B_M, assuming each woman proposes to at most half* *the men in B_M, no man in B_M receives more than $e_1 \ln n$ of these proposals, with failure probability $n \cdot \exp(-4d \ln n(\mu_u/\mu_\ell)/3)$, where $e_1 = 8d(\mu_u/\mu_\ell)$.*

Proof. Let m be a man in B_M. First, we bound the probability that a proposal is to man m. By Lemma 2.43, the ratio of probabilities for the proposals to men in B_M is bounded by $2\mu_u/\mu_\ell$, and by assumption, as least $\frac{1}{2}|B_M|$ have not yet been proposed to. Therefore, the probability that a proposal is to man m is at most $2(\mu_u/\mu_\ell) \cdot (2/|B_M|)$. Thus the expected number of woman-low proposals m receives is at most

$$\frac{2\mu_u}{\mu_\ell} \cdot \frac{2}{|B_M|} \cdot d \ln n|B_M| \leqslant \frac{4\mu_u}{\mu_\ell} \cdot d \ln n.$$

The upper bounds on these probabilities are based on the women's private scores for m, and therefore we can use stochastic dominance to justify applying a Chernoff bound. Thus, the number of these proposals is at most $8d(\mu_u/\mu_\ell) \ln n$ with failure probability at most $\exp(-4d \ln n(\mu_u/\mu_\ell)/3)$. A union bound over the men in B_M yields the final result. □

Let $B_{M,h}$ denote the set of men in B_M who eventually receive a man-high proposal, and $B_{M,\ell}$ denote the set $B_M \setminus B_{M,h}$.

Lemma 2.46. *If a woman $w_i \in B_W$ is currently matched with a man m_j in $B_{M,\ell}$, the probability that the next woman-low and man-low proposal is to m_j is at most $4(\mu_u/\mu_\ell)/|B_M|$.*

Proof. By Lemma 2.43, the ratio of probabilities for the proposals to men in B_M is bounded by $2\mu_u/\mu_\ell$, and by assumption, as least $\frac{1}{2}|B_M|$ men have not yet been proposed to. Therefore, the probability that a proposal is to m_j is at most $2(\mu_u/\mu_\ell) \cdot (2/|B_M|)$. □

Lemma 2.47. *Suppose that $8d(\mu_u/\mu_\ell) \ln n$ is an integer. Let $w \in B_W$. If w has at least $e_2(\ln n)^2$ man-low and woman-low edges to men in B_M, then the probability that she is unmatched after $d \ln n|B_M|$ man and women-low proposals to B_M is at most $\exp\left(-\frac{8}{3}d \cdot e_1\left(\frac{\mu_u}{\mu_\ell}\right)^2 \cdot (\ln n)^2\right)$, where $e_2 = 16d \cdot e_1\left(\frac{\mu_u}{\mu_\ell}\right)^2$.*

Proof. Suppose w is currently matched to a man in $B_{M,\ell}$. By Lemma 2.46, the probability that she is bumped (i.e. loses her current match) by the next man and women-low proposal to B_M is at most $4(\mu_u/\mu_\ell)/|B_M|$.

Therefore, over the course of $d \ln n|B_M|$ such proposals, she is bumped at most an expected $4d(\mu_u/\mu_\ell)\ln n$ times. Using stochastic dominance, we can apply a Chernoff bound, which shows she is bumped at most $8d(\mu_u/\mu_\ell)\ln n - 1$ times with failure probability at most $\exp\left(-\frac{4}{3}d(\mu_u/\mu_\ell)\ln n\right)$.

We now bound the probability that w tentatively matches with that man. By Lemma 2.45, m receives at most $e_1 \ln n$ proposals (including the current proposal by w). Each proposal has probability at most $\mu_u\Delta$ and at least $\mu_\ell\Delta$ of being in a Δ range of loss for the man, and therefore w's proposal produces the least loss among these up to $e_1 \ln n$ proposals with probability at least

$$\frac{\mu_\ell}{(e_1 \ln n - 1)\cdot \mu_u + \mu_\ell} \geqslant \frac{1}{e_1 \ln n}\cdot\frac{\mu_\ell}{\mu_u}.$$

Note that the bounds for each man are independent as they depend on the private scores of that man for the proposals he has received.

Therefore, to end up matched after these $d \ln n|B_M|$ proposals, it suffices that w make an expected

$$8d\frac{\mu_u}{\mu_\ell}\ln n \cdot e_1 \ln n \cdot \frac{\mu_u}{\mu_\ell} = 8d \cdot e_1\left(\frac{\mu_u}{\mu_\ell}\right)^2 \cdot (\ln n)^2 \ \text{ proposals.} \tag{2.5}$$

Then, by a Chernoff bound, she makes at most $16d \cdot e_1\left(\frac{\mu_u}{\mu_\ell}\right)^2 \cdot (\ln n)^2$ proposals with failure probability at most $\exp\left(-\frac{8}{3}d \cdot e_1\left(\frac{\mu_u}{\mu_\ell}\right)^2 \cdot (\ln n)^2\right)$. □

Lemma 2.48. *Each woman in B_W has at least $e_2(\ln n)^2$ man and woman-low edges to B_M with failure probability at most $\exp\left(-e_2 \ln^2 n/3\right)$, if $v \leqslant \frac{1}{8\eta^2}\left(\frac{\mu_\ell}{\rho_\ell}\right)^2\left(\frac{\mu_u}{\mu_\ell}\right)^5$, $t_m \geqslant 2$, and $t_w \geqslant 2$.*

Proof. As in the proof of Lemma 2.42, the probability that an edge from a woman $w_j \in B_W$ to a man $m_i \in B_M$ is man-low is at least $\frac{3}{4}\overline{\sigma} t_w^2 \rho_\ell$ and this depends on the man's private score for this edge; similarly the probability that it is woman low is at least $\frac{3}{4}\overline{\sigma} t_m^2 \rho_\ell$ and depends on the woman's private score for the edge. As $\mathcal{B}_8$ does not occur, by Lemma 2.38, $|B_M| \geqslant \frac{1}{2}\sigma_m \cdot n$. Thus, the expected number of man and woman-low edges from w_j to B_M is at least

$$\begin{aligned}\frac{3}{4}\overline{\sigma} t_w^2 \rho_\ell \cdot \frac{3}{4}\overline{\sigma} t_m^2 \rho_\ell \cdot \tfrac{1}{2}\sigma_m \cdot n &\geqslant \frac{1}{4} \cdot \frac{\overline{\sigma}^6 \cdot n}{\sigma_w^2 \sigma_m} (\rho_\ell)^2 \\ &\geqslant 32 \cdot 128 (c+2)^2 \left(\frac{\mu_\ell}{\rho_\ell}\right)^2 \frac{(\ln n)^2}{\eta^2 \nu} \\ &\geqslant 2e_2 (\ln n)^2, \quad \text{if}\end{aligned}$$

$$\begin{aligned}2e_2 = 16 d e_1 \left(\frac{\mu_u}{\mu_\ell}\right)^2 &= 8 \cdot 16 \left[2(c+2)\left(\frac{\mu_u}{\mu_\ell}\right)\right]^2 \cdot \left(\frac{\mu_u}{\mu_\ell}\right)^3 \\ &\leqslant 32 \cdot 128 (c+2)^2 \left(\frac{\mu_\ell}{\rho_\ell}\right)^2 \frac{1}{\eta^2 \nu}\end{aligned}$$
$$\text{i.e., if } 8\left(\frac{\mu_\ell}{\rho_\ell}\right)^2 \geqslant \left(\frac{\mu_\ell}{\mu_u}\right)^5 \eta^2 \nu.$$

Applying stochastic dominance, by a Chernoff bound, the number of edges is at least $e_2 \ln^2 n$ with probability at most $\exp(-e_2 \ln^2 n/3)$. □

Lemma 2.49. *All the women in T_W are matched with failure probability at most $O(n^{-(c+1)})$.*

Proof. Let w_j be a woman in T_W. If $r_{m_j} \geqslant \overline{\sigma}$, then the truncation does not remove any of the acceptable edges to w_j and so the previous analysis shows w_j is matched with failure probability $O(n^{-(c+2)})$.

So now suppose that $r_{m_j} < \overline{\sigma}$. Consider a run of man-proposing DA with the edge set cut at w_j and $\frac{1}{4} r_{w_j}$. Now, the acceptable edges are all woman-high. Furthermore, the acceptable edges cause a man m_i a loss of at most $U(r_{w_i}, 1) - U(r_{w_i} - \overline{\sigma} t_w^2, 1)$, and these are edges that are not truncated by m_i. The proof of Theorem 2.17 shows that such a woman m_j is matched using these edges with failure probability $O(n^{-(c+2)})$. □

Lemma 2.50. *If $\mathcal{B}_8$ does not occur, then, for large enough n, Assumption 1 holds with failure probability* $\exp(-|B_M|/8)$.

Proof. Assumption (i) holds if $\overline{\sigma}t_w^2\rho_u \leqslant \frac{1}{2}$, i.e. if $\overline{\sigma}^3/\sigma_w^2\rho_u \leqslant \frac{1}{2}$, i.e. if $128(c+2)\ln n/(\rho_\ell^2 n)\cdot n^{2/3}/\eta^2 \cdot \rho_u \leqslant \frac{1}{2}$; this holds if $(n^{1/3}/\ln n) \geqslant 256(c+2)\rho_u/[(\rho_\ell^2\eta^2]$, which is true for large enough n.

Assumption (ii) holds if each woman in B_W has at most $\frac{1}{2}|B_M|$ untruncated edges to men in B_M. The probability that an edge (m_i, w_j) is not truncated by m_i is at most $(r_{w_j}+\overline{\sigma}t_m^2)\rho_u \leqslant 2\overline{\sigma}t_m^2\rho_u$, and the probability that it is not truncated by w_j is at most $(r_{m_i}+\overline{\sigma}t_w^2)\rho_u \leqslant 2\overline{\sigma}t_w^2\rho_u$. Thus the expected number of untruncated edges from a woman $w \in B_W$ to the men in B_M is at most

$$2\overline{\sigma}t_w^2\rho_u \cdot 2\overline{\sigma}t_m^2\rho_u \cdot |B_M| \leqslant 4\frac{\overline{\sigma}^6}{\sigma_w^2\sigma_m^2}\rho_u^2|B_M| \leqslant \frac{(256(c+2)\ln n\rho_u)^2}{\rho_\ell^4\eta^2\nu^2 n^{2/3}}|B_M| \leqslant \tfrac{1}{4}|B_M|,$$

if n is large enough.

Note that the bounds on the probabilities are due to the men's and women's independent private scores for these edges. Thus, using stochastic dominance, by means of a Chernoff bound, we obtain that the number of these edges is at most $\frac{1}{2}|B_M|$ with failure probability $\exp(-|B_M|/8)$. □

Lemma 2.51. *The run of woman-proposing DA with the truncated edge sets matches every woman (and man) with failure probability n^{-c} if n is large enough, if $\nu = \frac{64}{\eta^2}\cdot\left(\frac{1}{\rho_\ell}\right)^2\cdot\left(\frac{\mu_\ell}{\mu_u}\right)^4$ and η satisfies* $4\eta^3\rho_\ell^4(\mu_u/\mu_\ell)^4 \cdot \exp\left(-\eta^3\rho_\ell^2/128\right) \leqslant \frac{1}{10}$.

Proof. As any unmatched woman in B_W will keep proposing until she runs out of proposals, we deduce from Lemmas 2.47 and 2.48 that all the women in B_W are matched, modulo the lemma's failure probability. By Lemma 2.49, all the women in T_W are matched, modulo the lemma's failure probability. Thus all the women are matched.

This entails the following constraints, from Lemmas 2.40, 2.41, 2.42, Corollary 3, Lemmas 2.45, 2.47, 2.48, 2.41, respectively.

$$
\begin{aligned}
n^{2/3} &\geqslant 3 \\
\eta &\geqslant 6\nu \\
t_m &\geqslant 2 \\
d &= 2(c+2)\mu_u/\mu_\ell \\
e_1 &= 16d(\mu_u/\mu_\ell) = 8(c+2)(\mu_u/\mu_\ell)^2 \\
e_2 &= 8d \cdot e_1(\mu_u/\mu_\ell)^2 = 8(c+2)^2(\mu_u/\mu_\ell)^4 \\
t_w &\geqslant 2 \\
\nu &\leqslant \frac{1}{8\eta^2} \cdot \left(\frac{\mu_\ell}{\rho_\ell}\right)^2 \cdot \left(\frac{\mu_u}{\mu_\ell}\right)^5 \\
\tfrac{1}{10} &\geqslant 4(\eta/\nu) \cdot \exp\bigl(-(\eta/2)^3\rho_\ell^2/128\bigr)
\end{aligned}
$$

We set $\nu = \frac{1}{8\eta^2} \cdot \left(\frac{\mu_\ell}{\rho_\ell}\right)^2 \cdot \left(\frac{\mu_u}{\mu_\ell}\right)^5$. The final constraint becomes

$$\eta^3(\rho_\ell/\mu_\ell)^2(\mu_\ell/\mu_u)^5 \cdot \exp\bigl(-(\eta/2)^3\rho_\ell^2/128\bigr) \leqslant \tfrac{8}{10}.$$

In addition, we need to satisfy $\eta \geqslant 6\nu$. Clearly, $\eta = O(1)$ suffices.

Finally, to ensure $t_w \geqslant 2$ it suffices to have

$$\left(\frac{128(c+2)\ln n}{\rho_\ell^2}\right)^{1/3} \geqslant 2\eta,$$

and clearly this holds if n is large enough. As $t_m > t_w$, this also ensures that $t_w \geqslant 2$.

We also assume that $8d(\mu_u/\mu_\ell)\ln n = 8(c+2)(\mu_u/\mu_\ell)^2 \ln n$ is an integer (in Lemma 2.47). This can be achieved by increasing μ_u slightly.

The overall failure probability obtained by summing the terms in Lemmas 2.48, 2.47, 2.45, 2.44, 2.41, 2.40, 2.39, 2.38, 2.49, 2.50, and Corollary 3, plus ruling out $\mathcal{B}_1$–$\mathcal{B}_3$, is at most

$$\exp\big(-e_2(\ln n)^2/3\big) + \exp\big(-\tfrac{8}{3}d \cdot e_1(\mu_u/\mu_\ell)^2 \cdot (\ln n)^2\big) + n \cdot \exp(-4d \ln n(\mu_u/\mu_\ell)/3)$$
$$+ \exp\big(-d\eta v^2 \rho_u \rho_\ell |B_M|/[48(c+2)]\big) + \exp(-|B_M|/24) + 2\exp\big(-\eta \cdot n^{2/3}/3\big)$$
$$+ \exp(-\sigma_m \cdot n) + 6\exp(-\sigma_m \cdot n/8) + \exp(-|B_M|/8) + O(n^{-(c+1)}).$$

This totals $O(n^{-(c+1)})$, which is bounded by n^{-c} for large enough n. □

Proof. (of Theorem 2.18) Lemma 2.51 shows that, with probability at least $1 - n^{-c}$, there exists a stable matching, in which every man and woman obtains a match with a loss of less than $L^m_{t_m}$ and $L^w_{t_w}$, respectively; it results from the men with public rating $\overline{\sigma}t$ implementing reservation strategies with reservation thresholds L^m_t, for $t < t_m$, and the remaining men using the reservation threshold $L^m_{t_m}$. The edges meeting this constraint are the acceptable edges for this run of DA. By Theorem 2.15, w.h.p, no man m gets utility greater than $U(r^m, 1) + \Theta([\ln n/n]^{1/3})$, and an analogous bound applies to the women. Thus, the most a man could gain by deviating from the equilibrium strategy, in terms of his expected utility, is

$$n^{-c} \cdot 2 + (1 - n^{-c}) \cdot (\Theta([\ln n/n]^{1/3}) + L^m_{t_m}).$$

Since $L^m_{t_m} = \Theta(\ln n/n^{1/3})$, this is an ε-Bayes-Nash equilibrium with $\varepsilon = \Theta(\ln n/n^{1/3})$.

Further notice that, for each agent, the number of acceptable edges is at most $\Theta(\ln^2 n)$; furthermore, this bound improves to at most $\Theta(\ln n)$ for all agents outside the bottom $\Theta([\ln n/n]^{1/3})$ fraction of agents. □

2.12 Numerical Simulations

We present several simulation results which are complementary to our theoretical results. Throughout this section, we focus on the linear separable model.

2.12.1 NRMP Data

We used NRMP data to motivate some of our choices of parameters for our simulations. The NRMP provides extensive summary data [nrmp.org 2021]. We begin by discussing this data.

Over time, the number of positions and applicants has been growing. We mention some numbers for 2021. There were over 38,000 positions available and a little over 42,000 applicants. The main match using the DA algorithm (modified to allow for couples, who comprise a little over 5% of the applicants) filled about 95% of the available positions. The NRMP also ran an aftermarket, called SOAP, after which about 0.5% of the positions remained unfilled.

The positions cover many different specialities. These specialities vary hugely in the number of positions available, with the top 11, all of size at least 1,000, accounting for 75% of the positions. In addition, about 75% of the doctors apply to only one speciality. We think that as a first approximation, w.r.t. the model we are using, it is reasonable to view each speciality as a separate market. Accordingly, we have focused our simulations on markets with 1,000–2,000 positions (though the largest speciality in the NRMP data had over 9,000 positions).

On average, doctors listed 12.5 programs in their preference lists, hospital programs listed 88 doctors, and the average program size was 6.5 (all numbers are approximate). While there is no detailed breakdown of the first two numbers, it is clear they vary considerably over the individual doctors and hospitals. For our many-to-one simulations we chose to use a fixed size for the hospital programs. Our simulations cause the other two numbers to vary over the individual doctors and programs because the public ratings and private scores are chosen by a random process.

2.12.2 Numbers of Available Edges

The first question we want to answer is how long do the preference lists need to be in order to have a high probability of including all acceptable edges, for all but the bottommost agents?

We chose bottommost to mean the bottom 20% of the agents, based on where the needed length of the preference lists started to increase in our experiments for $n = 1{,}000-2{,}000$.

We ran experiments with $\lambda = 0.5, 0.67, 0.8$, corresponding to the public rating having respectively equal, twice, and four times the weight of the private scores in their contribution to the utility. We report the results for $\lambda = 0.8$. The edge sets were larger for smaller values of λ, but the results were qualitatively the same. We generated 100 random markets and determined the smallest value of L that ensured all agents were matched in all 100 markets. $L = 0.12$ sufficed. In Figure 2.2, we show results by decile of women's rank (top 10%, second 10%, etc.), specifically the average length of the preference list and the average number of edges proposed by a woman in woman-proposing DA, over these 100 randomly generated markets. We also show the max and min values over the 100 runs; these can be quite far from the average value. Note that the min values in Figure 2.2(a) are close to the max values in Figure 2.2(b), which suggests that being on the proposing side does not significantly reduce the value of L that the women could use compared to the value the men use. We also show data for a typical single run in Figure 2.3.

We repeated the simulation for the many-to-one setting. In Figure 2.4, we show the results for 2000 workers and 250 companies, each with 8 positions. Now, on average, a typical worker (i.e. among the top 80%) has an average preference list length of 55 and makes 7 proposals.

The one-to-one results show that for non-bottommost agents, the preference lists have length 150 on the average, while women make 30 proposals on the average (these numbers are slightly approximate). What is going on? We believe that the most common matches provide a small loss or gain ($\Theta(n^{-1/3})$ in our theoretical bounds) as opposed to the maximum loss possible ($\Theta(n^{-1/3} \ln^{1/3} n)$ in our theoretical bounds), as is indicated by our distribution bound on the losses

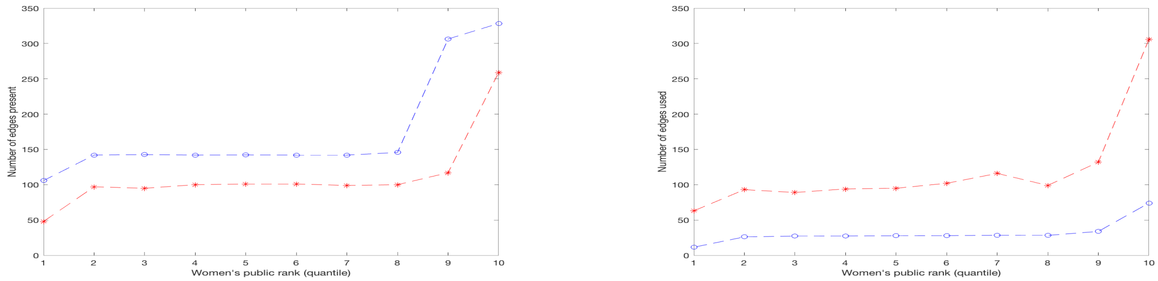

(a) Number of edges in the acceptable edge set, per woman, by decile; average in blue with circles, minimum in red with stars. ($n = 2{,}000$, $\lambda = 0.8$, $L = 0.12$.)

(b) Number of edges in the acceptable edge set proposed during the run of DA, per women, by decile; average in blue with circles, maximum in red with stars.

Figure 2.2: One-to-one case ($n = 2000$): summary statistics.

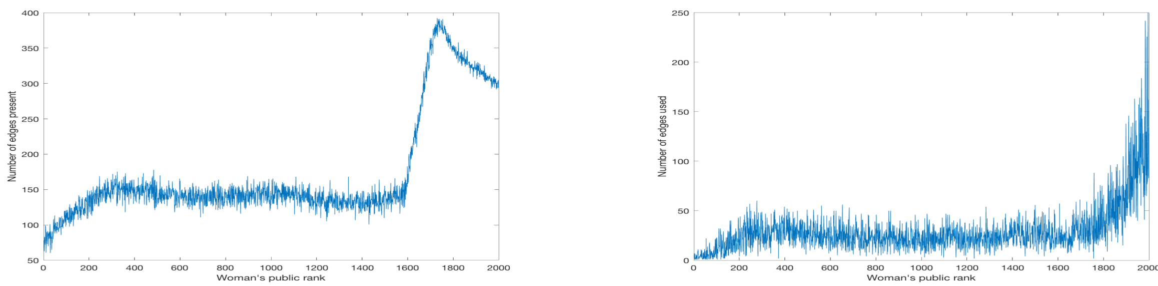

(a) Number of edges in the acceptable edge set for each woman.

(b) Number of edges in the acceptable edge set proposed by each woman.

Figure 2.3: One-to-one case ($n = 2000$): a typical run.

(see item 4 in Section 2.8.1). The question then is where do these edges occur in the preference list, and the answer is about one fifth of the way through (for one first has the edges providing a gain, which only go to higher up agents on the opposite side, and then one has the edges providing a loss, and these go both up and down). However, a few of the women will need to go through most of their list, as indicated by the fact that the max and min lines (for example in Figure 2.4) roughly coincide.

This effect can also be seen in the many-to-one experiment but it is even more stark on the worker's side. The reason is that the number of companies with whom a worker w might match which are above w, based on their public ratings alone, is $\Theta(L_c n_c)$, while the number below w

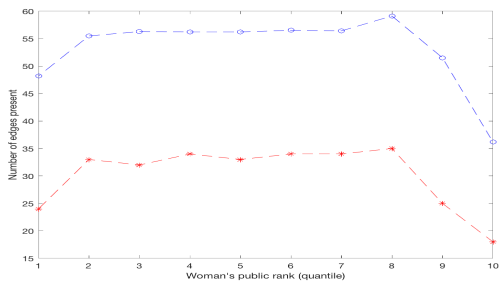

(a) Many to One Setting: Number of edges in the acceptable edge set per worker, by decile; average in blue with circles, minimum in red with stars. (n_w = 2,000, $d = 8$, $\lambda = 0.8$, $L_c = 0.14$, $L_w = 0.24$.)

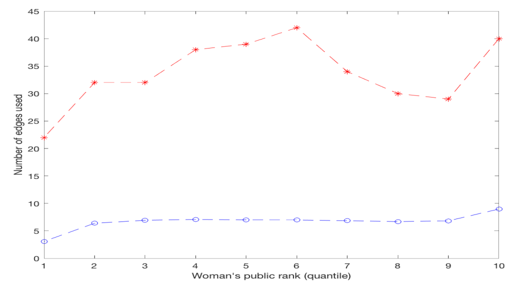

(b) Number of edges in the acceptable edge set proposed during the run of DA, per worker, by decile; average in blue with circles, maximum in red with stars.

Figure 2.4: Many to One Setting ($n = 2000$).

is $\Theta(L_w n_c)$, a noticeably larger number. (See Section 2.9.1 for a proof of these bounds.) The net effect is that there are few edges that provide w a gain, and so the low-loss edges, which are the typical matches, are reached even sooner in this setting.

Now we turn to why the number of edges in the available edge set per woman changes at the ends of the range. There are two factors at work. The first factor is due to an increasing loss bound as we move toward the bottommost women, which increases the sizes of their available edge sets. The second factor is due to public ratings. For a woman w the range of men's public ratings for its acceptable edges is $[r_m - \Theta(\overline{L}), r_m + \Theta(\overline{L})]$, where m is aligned with w. But at the ends a portion of this range will be cut off, reducing the number of acceptable edges, with the effect more pronounced for low public ratings. Because $\lambda = 0.8$, initially, as we move to lower ranked women, the gain due to increasing the loss bound dominates the loss due to a reduced public rating range, but eventually this reverses. Both effects can be clearly seen in Figure 2.3(a), for example.

2.12.3 Unique Stable Partners

Another interesting aspect of our simulations is that they showed that most agents have a unique stable partner. This is similar to the situation in the popularity model when there are short preference lists, but here this result appears to hold with full length preference lists. In Figure 2.5, we show the outcome on a typical run and averaged over 100 runs, for $n = 2{,}000$ in the one-to-one setting. We report the results for the men, but as the setting is symmetric they will be similar for the women. On the average, among the top 90% of agents by rank, 0.5% (10 of 1,800) had more than one stable partner, and among the remainder another 2% had multiple stable partners (40 of 200).

Also, as suggested by the single run illustrated in Figure 2.5(a), the pair around public rank 1,600 and the triple between 1,200 and 1,400 have multiple stable partners which they can swap (or exchange via a small cycle of swaps) to switch between different stable matchings. This pattern is typical for the very few men with multiple stable partners outside the bottommost region.

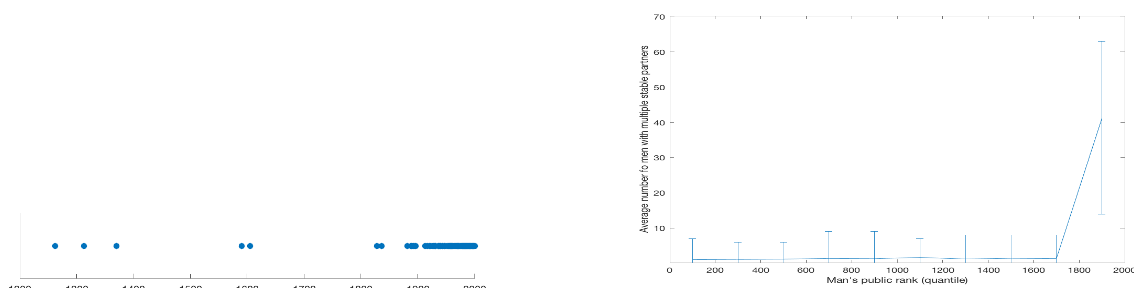

(a) Public ranks of men with multiple stable partners in a typical run.

(b) Average numbers of men with multiple stable partners, by decile.

Figure 2.5: Unique stable partners, one-to-one setting ($n = 2000$).

2.12.4 Constant Number of Proposals

Our many-to-one experiments suggest that the length of the preference lists needed by our model are larger than those observed in the NRMP data. In addition, even though there is a simple rule for identifying these edges, in practice the communication that would be needed to identify these edges may well be excessive. In light of this it is interesting to investigate what can be done when the agents have shorter preference lists.

We simulated a strategy where the workers' preference lists contain only a constant number of edges. We construct an *Interview Edge Set* which contains the edges (w, c) satisfying the following conditions:

1. Let r_w and r_c be the public ratings of w and c respectively. Then $|r_w - r_c| \leqslant p$.
2. The private score w has for c as well as the private score of c for w are both greater than q.

We choose the parameters p and q so as to have 15 edges per agent on average. Many combinations of p and q would work. We chose a pair that caused relatively few mismatches. We then ran worker proposing DA on the Interview Edge Set.

One way of identifying these edges is with the following communication protocol: the workers signal the companies which meet their criteria (the workers' criteria); the companies then reply to those workers who meet their criteria. In practice this would be a lot of communication on the workers's side, and therefore it may be that an unbalanced protocol where the workers use a larger q_w as their private score cutoff and the companies a correspondingly smaller q_c is more plausible. Clearly this will affect the losses each side incurs when there is a match, but we think it will have no effect on the non-match probability, and as non-matches are the main source of losses, we believe our simulation is indicative. We ran the above experiment with $p = 0.19$ and $q = 0.60$, with the company capacity being 8. Figure 2.6(a) shows the locations of unmatched workers in a typical run of this experiment while 2.6(b) shows the average numbers of unmatched workers per quantile (of public ratings) over 100 runs. We observe that the number of unmatched workers is very low (about 1.5% of the workers) and most of these are at the bottom of the public rating range.

Figure 2.6(c) compares the utility obtained by the workers in the match obtained by running worker-proposing DA on the Interview Edge Set to the utility they obtain in the worker-optimal stable match. We observe that only a small number of workers have a significantly worse outcome when restricted to the Interview Edge Set.

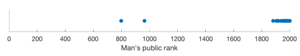

(a) Public ranks of unmatched workers in a typical run.

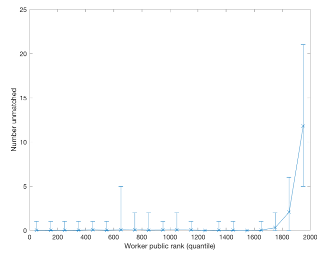

(b) Average numbers of unmatched workers by public rating decile.

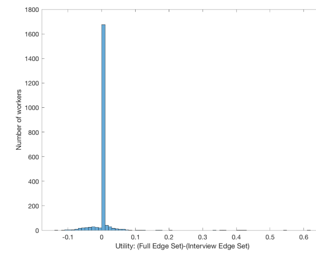

(c) Distribution of workers' utilities with worker-proposing DA: (full edge set result) −(Interview edge set result)

Figure 2.6: Constant number of proposals ($n = 2000$).

2.13 Additional Numerical Simulations

Here we provide another set of the experiments, but for $n = 1{,}000$ instead of 2,000. The relative weight of public ratings and private scores is unchanged ($\lambda = 0.8$).

2.13.1 Numbers of Available Edges

2.13.1.1 One-to-one

$n = 1{,}000$, $\lambda = 0.8$, $L = 0.15$, 100 runs.

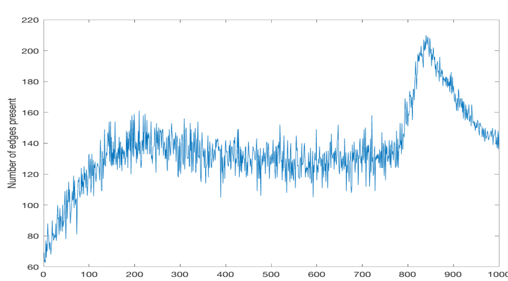

(a) Number of edges in the acceptable edge set for each woman.

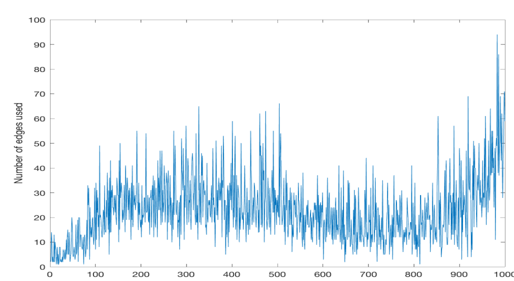

(b) Number of edges in the acceptable edge set proposed by each woman.

Figure 2.7: One-to-one case: Outcome in a typical run ($n = 1000$).

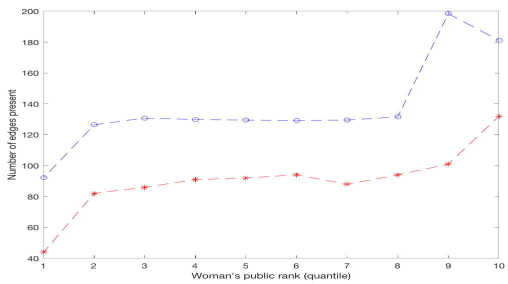

(a) Number of edges in the acceptable edge set, per woman, by decile; average in blue with circles, minimum in red with stars.

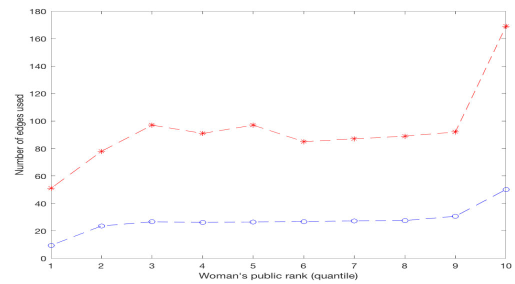

(b) Number of edges in the acceptable edge set proposed during the run of DA, per women, by decile; average in blue with circles, maximum in red with stars.

Figure 2.8: One-to-one case ($n = 1000$): summary statistics.

2.13.1.2 Many-to-one

$n = 1{,}000$, $\lambda = 0.8$, $d = 4$, $L_c = 0.16$, $L_w = 0.25$, 100 runs.

We chose to present the results for $d = 4$ rather than 8 (as used in the $n = 2{,}000$ experiments) because the needed value for L_w with $d = 8$ leads to very large acceptable edge sets, which we do not consider an interesting case.

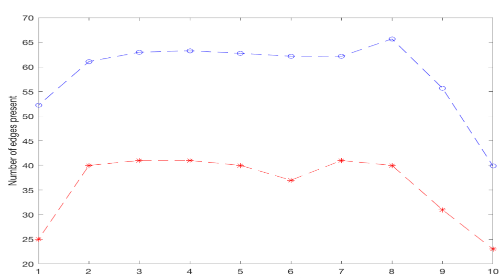

(a) Number of edges in the acceptable edge set, per woman, by decile; minimum in red with stars, average in blue with circles. ($n_w = 1{,}000$, $d = 4$, $\lambda = 0.8$, $L_c = 0.15$, $L_w = 0.25$.)

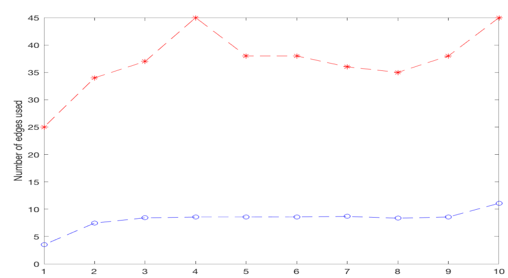

(b) Number of edges in the acceptable edge set proposed during the run of DA, per woman, by decile; maximum in red with stars, average in blue with circles.

Figure 2.9: Many to One Setting ($n = 1000$)

2.13.2 Unique Stable Partners

100 runs; 38 men have multiple stable partners in the typical run shown.

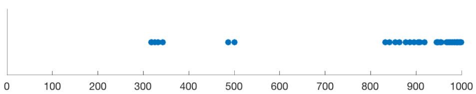

(a) Public rank of men with multiple stable partners in a typical run.

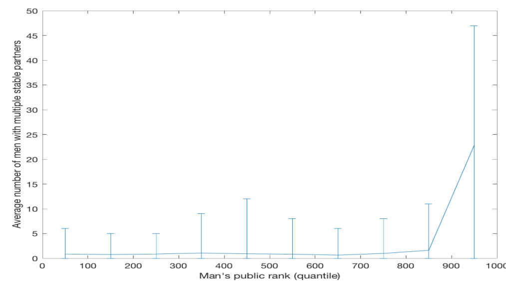

(b) Average numbers of men with multiple stable partners, by decile.

Figure 2.10: Unique stable partners, one-to-one setting ($n = 1000$).

2.13.3 Constant Number of Proposals

$r = 0.19$, $q = 0.60$, company capacity = 4, 100 runs.

(a) Public ranks of unmatched workers in a typical run.

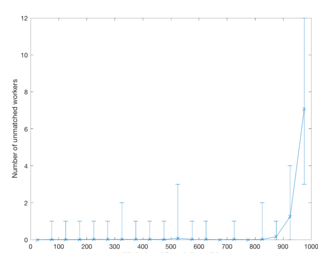

(b) Average number of unmatched workers, by decile.

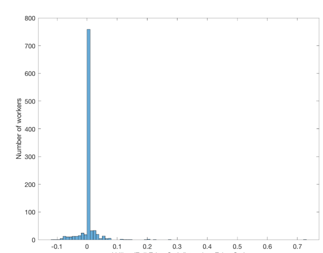

(c) Distribution of workers' utilities with worker-proposing DA: (full edge set result) −(Interview edge set result)

Figure 2.11: Constant number of proposals ($n = 1000$).

2.14 Discussion and Open Problems

Our work shows that in the bounded derivatives model, apart from a sub-constant fraction of the agents, each of the other agents has $O(\ln n)$ easily identified edges on their preference list which cover all their stable matches w.h.p.

As described in Section 2.12, our experiments for the one-to-one setting yield a need for what appear to be impractically large preference lists. While the results in the many-to-one setting are more promising, even here the preference lists appear to be on the large side. Also, while our rule for identifying the edges to include is simple, in practice it may well require too much communication to identify these edges. At the same time, our outcome is better than what is achieved in practice: we obtain a complete match with high probability, whereas in the NRMP setting a small but significant percentage of positions are left unfilled. Our conclusion is that it remains important to understand how to effectively select smaller sets of edges.

In the popularity model, it is reasonable for each agent to simply select their favorite partners. But in the current setting, which we consider to be more realistic, it would be an ineffective strategy, as it would result in most agents remaining unmatched. Consequently, we believe the main open issue is to characterize what happens when the number of edges k that an agent can list is smaller than the size of the allowable edge set. We conjecture that following a simple protocol for selecting edges to list, such as the one we use in our experiments (see Section 2.12.4), will lead to an ε-Bayes-Nash equilibrium, where ε is a decreasing function of k. Strictly speaking, as the identification of allowable edges requires communication, we need to consider the possibility of strategic communication, and so one would need to define a notion of ε-equilibrium akin to a Subgame Perfect equilibrium. We conjecture that even with this, it would still be an ε-equilibrium.

Finally, it would be interesting to resolve whether the experimentally observed near uniqueness of the stable matching for non-bottom agents is a property of the linear separable model. We conjecture that in fact it also holds in the bounded derivatives model.

3 Selecting a Match: Exploration vs Decision

3.1 Introduction

What strategies make sense when deciding whether to commit to a long-term relationship? We are interested in pairings between members of two sets of agents, such as an employer offering a job and a worker accepting, a woman (or man) proposing marriage to a person of the opposite sex,[1] a landlord agreeing to rent an apartment to a potential renter.

The key feature of these relationships is that the longer they last, the greater the utility they provide; for simplicity, we assume this utility is linear in the duration of the match. Nonetheless, as a rule agents do not choose to match as soon as they receive a proposal, for different potential partners may provide different utilities. An employer may be supportive or not, a marriage may be happy or not; the possibilities are myriad. Agents seek to assess the utility of a proposed match and then decide whether to accept or keep searching (such an assessment might be implicit). These judgements can be based on some combination of idiosyncratic factors and commonly shared perspectives. Both sides of a potential match are making this assessment, and a match happens only if both sides accept it.

[1] Single-sex marriages could also be studied, but then there would be just one set of agents. In fact, this does not appear to significantly affect our results, but in this work we have focused on the case of two sets of agents.

Assessing potential matches takes time and therefore an agent can consider only a relatively small number of potential matches at any one time. In many circumstances, choices are offered on a take it or lose it basis. Typically, job offers are made with a short decision window. While marriage or its equivalents have many cultural variations, as a rule offers of marriage when made are accepted or declined; it would be unusual to collect multiple offers and only then decide (in the somewhat unlikely event the parties on the other side would be willing to wait). Again, for simplicity, we assume agents can consider only one match at a time.

Furthermore, agents are aware of time slipping by. An unemployed worker cannot afford to stay unemployed indefinitely. Businesses wish to fill open positions promptly as they need workers to carry out the duties of these open positions. Many men and women appear to want to pair sooner rather than later (whether the pairing is called marriage or not). We see two forces at work here: one is the ongoing utility from a match, which starts only when the match is formed. The second is that the longer an agent waits the shorter the duration of the match they can offer.

We are interested in two questions:

What decision rule make sense and how can their effectiveness be measured?

Each potential decision rule provides a balance between the urge to form a match soon so as to have a longer time in which to enjoy it, and the desire to continue searching in the hopes of finding a better match.

The equilibrium properties of decision rules have been studied previously in models with a continuum population, a continuum model for short [Adachi 2003; Burdett and Coles 1997, 1999; Burdett and Wright 1998; Smith 2006; Shimer and Smith 2000; Bloch and Ryder 2000; Eeckhout 1999; Lauermann and Nöldeke 2014; McNamara and Collins 1990; Damiano et al. 2005]. In these works, agents are assumed to arrive according to a variety of processes, such as a Poisson process. In some of these works, they are also assumed to use time discounting of future utility. Either they have infinite lifetimes in which to seek matches or they depart—die—according to another process. We discuss this in more detail in the related work section below. Each agent has an intrinsic appeal, a numeric value, called *charm* in Burdett and Coles [Burdett and Coles 1999]. The utility an agent derives from a match is assumed to be an increasing function of their partner's charm. Agents receive match proposals at a fixed rate and agents either accept or reject a match immediately; for a match to succeed both participating agents must agree to it. One natural class of agent strategies are reservation strategies; an agent will accept a proposed match exactly if the partner has charm at least c. Typically the chosen c is a function of the agent's own charm. The right choices of reservations c yield equilibrium strategies.

In contrast, we study this problem in a discrete, albeit stochastic, setting. By this we mean that a finite number of agents arrive at each time step; we also choose time to be discrete. In addition, we model lifetimes differently, viewing all lives as having duration T. This has the effect of making agents less demanding over time, which we believe is a real effect, and an effect that will not arise with a departure rate that stays the same over time.

Discreteness introduces variance, which leads to localized imbalances in the numbers of men and women (by localized, we mean agents of a given age and charm). The analysis and bounding of these imbalances are the largest challenge we face, and while asymptotically small, for moderate values of our parameters these are non-trivial quantities, as confirmed by our simulation results. This is in sharp contrast to a continuum setting, where there will be no variance. Finally, it is not clear that our setting will converge to an equilibrium or near-equilibrium, and while our simulations for moderate parameter values suggest a certain level of stability, they also show that there is continuing substantial variability. In any event, our concern is to understand the quality of the outcomes: in a sense we make precise shortly, our model achieves near-optimal utility with high probability.

Related Work Rogerson [Rogerson et al. 2005] surveyed issues of search cost and bargaining in job markets. More recently, Chade, Eeckout and Smith [Chade et al. 2017] gave a broad survey of matching in economic models, covering search with and without costs, and settings with and without transferable utility. We focus on settings with search costs and no transferable utility. Even in this domain there are many works. We characterize these works w.r.t. multiple dimensions.

The first is the treatment of time, both as regards arrivals and departures. Most papers assume agents remain in the market till they are matched. A few allow matches to be broken via a Poisson process (e.g., jobs end, partners divorce) and then the agents return to the market; see Shimer and Smith [Shimer and Smith 2000] and Smith [Smith 2006]. Others have agents ending their participation via various random processes: Burdett and Wright [Burdett and Wright 1998] use a Poisson process, Adachi [Adachi 2003] uses an exponential random variable, and Lauermann and Noldeke [Lauermann and Nöldeke 2014] use an exogeneous rate. Arrivals are similarly varied. Poisson processes are considered in Burdett and Coles [Burdett and Coles 1997], Smith [Smith 2006], and Shimer and Smith [Shimer and Smith 2000]. Other works consider cloning: when agents leave due to a match they are replaced by clones thereby keeping the available matches unchanged; see Adachi [Adachi 2003] and Burdett and Wright [Burdett and Wright 1998]. Fixed arrival rates: see Eeckhout [Eeckhout 1999], and Lauermann and Noldeke [Lauermann and Nöldeke 2014]. Finally, no new arrivals: see Damiano, Hao and Suen [Damiano et al. 2005], and McNamara and Collins [McNamara and Collins 1990].

The second dimension is the choice of utility model. These are all functions of the partner's charm, though there is considerable variation. The most common is that the utility an agent gains is a non-decreasing function, either linear [Burdett and Coles 1997] or more general [Smith 2006; Eeckhout 1999]; some papers allow for time discounting [Adachi 2003; Bloch and Ryder 2000]; the utility can be the product of the partners' charms [Damiano et al. 2005]; or it is given by independent random variables for each pair of agents [Burdett and Wright 1998; McNamara and Collins 1990]; another option is that the agents obtain their utility by dividing a reward which is a function of their individual charms [Shimer and Smith 2000].

The final dimension is the choice of equilibrium model. Most of the papers consider a steady state equilibrium; McNamara and Collins [McNamara and Collins 1990] consider Nash Equilibria, and Damiano, Hao and Suen [Damiano et al. 2005] analyze a multi-round dynamic equilibrium.

The tension between taking a choice now and waiting for potentially better options arises in multiple other domains, including secretary problems [Ferguson 1989], online matching [Karp et al. 1990], matching market thickening [Akbarpour et al. 2020; Baccara et al. 2020], and regret minimization [Blum and Mansour 2007]. In spirit, the secretary problem seems the most analogous as it involves a single decision, albeit by just a single agent. We discuss it briefly in the next paragraph. In contrast, online matching has a centralized decision maker that seeks to optimize the outcome of many choices. Regret minimization occurs in a distributed setting, however here each agent makes multiple decisions over time, with the goal of achieving a cumulatively good outcome; again, this seems quite distinct from our setting. Market thickening is used in contexts where a global matching is being computed, which seem unlike the random matches on offer in our setting.

The standard secretary problem is expressed in terms of ranks. A cardinal version was considered by Bearden [Bearden 2006]; here the goal is to maximize the expected value of the chosen secretary, with values uniform on $[0, 1]$. For each applicant the decision maker learns whether they are the best so far. Bearden shows the optimal strategy is to reject the first $\sqrt{n} - 1$ candidates, and then choose the first candidate who meets the "best so far" criteria. Clearly, the expected value of the selected secretary is $1 - \Theta(1/\sqrt{n})$, which is analogous to the bounds we obtain, although the settings appear quite distinct. Bearden argued that the payoff rule in this version of the problem is more natural that the classic version. The problem of maximizing the duration of a relatively best choice has also been considered [Ferguson 1989].

3.2 The Model

We consider a setting in which, at each time step, n agents enter a matching pool. Agents exit the pool either when they are matched or if they have been in the pool for T time steps. There are two types of agents, called men and women. Each match pairs a man with a women. At each time step the agents are paired uniformly at random. Each pair comprises a proposed match. Each agent in a pair can accept or reject the proposed match as they prefer; a match occurs only if both agents accept it.

In a discrete setting, a random pairing seems more natural than having pairs arrive one by one, for the process of pairing will proceed in parallel, and pairs are necessarily mutually exclusive. While in practice the pairings under consideration at any one time will not cover the whole of the smaller side of the population, considering a maximal matching seems a reasonable simplification.

We assume agents evaluate their potential partners using cardinal values, and furthermore these are common values: every agent of the opposite type (gender) has the same value v_i for agent i. In the terminology of Burdett and Coles, this is agent i's charm.

We associate two parameters v_i and t_i with agent i. v_i is the agent's charm and t_i is the total time remaining before agent i is forced to exit the pool. Agent i derives utility $v_j \cdot \min(t_i, t_j)$ when matched with agent j. We assume that the values lie in the range $[T, 2T)$, and that an agent's value, chosen when it enters the pool, is one of $\{T, T+1, \ldots, 2T-1\}$, picked uniformly at random. We note that the relative utilities of an agent are scale free; in other words, the range assumption is equivalent to assuming the values lie in the range $[1, 2)$. We could have used a separate discretization for the values, but we preferred to avoid an additional parameter. Furthermore, it would not affect the results qualitatively.

Entering agents are either male or female with equal probability.

Throughout this work it will be useful to view the market as a $T \times T$ size box, with agents located at grid points. The box is indexed by value on the horizontal axis and by time on the vertical axis. Consider the set of T points on the top edge: $\{(T, 0), (T + 1, 0), \ldots, (2T - 1, 0)\}$. Agents enter the market at one of these points, picked uniformly at random. At each time step, an agent either matches and leaves the box or moves down vertically by 1 unit. After T steps, if unmatched for all these times, the agent exits the box (at the bottom).

A Reasonable Notion of Loss In a single gender version of this setting, the total utility derived by the n agents that enter at any one time step is at most $n \cdot \sum_i v_i \cdot T$; in the two-gender case, by applying a Chernoff bound, one can obtain a similar bound with high probability. This bound can easily be achieved if all agents simply accept whatever match is proposed to them in the very first step in which they enter the matching pool. However such behavior seems implausible for high value agents, as their expected utility would be much smaller than what they might reasonably hope to achieve. Consequently, we set $v_i \cdot T$ as a reasonable target for i's achieved utility. Based on this, we define the *total loss* suffered by the agents to be:

$$\sum_i \left[v_i \cdot T - \text{utility obtained by agent } i\right]^+ .$$

This measure captures the intuition that agents who obtain less than their worth due either to a lower value partner, or to accepting a match only later on in the process, are suffering losses. We want to capture how much utility is lost compared to the benchmark in which each agent gets an equal value partner for the whole length T time period. It also addresses what is implausible about the naive solution, in which all agents immediately accept whatever match is proposed to them, and which maximizes the usual notion of social welfare.

It is not clear how to determine an optimal strategy, let alone whether it can be computed feasibly. For a truly optimal strategy would incorporate the effects of past variance, a level of knowledge that seems implausible in practice; and even an ex-ante optimal strategy seems out of reach. Instead, we will present a strategy, which we call the *reasonable strategy*, which seeks to ensure that if it is followed by all the players, then the total loss will be at most a constant factor larger than what could be achieved by the optimal strategy. Actually, we introduce two strategies, and the second one, called the *modified reasonable strategy*, is the one we analyze.

3.3 Our Results

We obtain a lower bound on the total loss suffered by agents; no matter their behavior, they will, with high probability, suffer an average loss of $\Omega(T\sqrt{T})$.

Theorem 3.1. *Suppose the matching market runs for τ time steps. If $16 \leq T+1 \leq n, c \geq 1, T \leq \tau \leq n^c$, and $n \geq 96T(2c+2)\ln n$, then, over τ time steps, whatever strategies the agents use, with probability at least $1-\frac{1}{4n^c}$, the average loss per agent is at least $\frac{T\sqrt{T}}{20}$.*

On the other hand, we construct a strategy profile, which if followed by all the agents, leads, with high probability, to a total loss of at most $O(T\sqrt{T})$.

Theorem 3.2. *Suppose $2T \leq \tau \leq n^c$, $c \geq 1$, $676 \leq T$, and $n \geq (3654+2436e^{12}+546(e^{12}+1)c)^2(3c+4)T^3(\log_2 n)^2\ln n$. Then, over τ time steps, if all agents follow the modified reasonable strategy, with probability at least $1-\frac{1}{n^c}$, the average loss per agent is at most $11T\sqrt{T}$.*

Our results hold for large n and T. Furthermore, Theorem 3.2 applies only when n is much larger than T. However, our numerical simulations suggest that similar results hold even for quite moderate values of n and T and also do not require n to be much bigger than T. To simplify the presentation, we assume that $T = 4^i$ for some integer $i > 0$, though the bounds extend to all values of T, possibly with somewhat larger constants.

3.4 Preliminaries

We review the notion of negative cylinder dependence and make a simple observation regarding the matching procedure.

Lemma 3.3. *Suppose there are m men and w women in total. Further suppose that for a given man x, there are w' women for which a proposed match would be accepted by both sides. Then a random match will provide man x such a match with probability $w'/\max\{m, w\}$.*

Proof. If there are at least as many women as men, every man will be offered a match, and the probability that it is accepted by both sides is w'/w. While if there are more men, a man will be offered a match with probability w/m, and thus the probability that he is offered an acceptable match is $w/m \cdot w'/w = w'/m$. □

Negative Dependence Consider a set of 0-1 valued valued random variables $\{X_i\}_{i=1}^n$. The set $\{X_i\}$ is *λ-correlated* if

$$E\Big[\prod_{i=1}^{n} X_i\Big] \leq \lambda \cdot \prod_{i=1}^{n} E\Big[X_i\Big],$$

where $\lambda \geq 1$. The set $\{X_i\}$ is *negative cylinder dependent* if $\{X_i\}$ and $\{1-X_i\}$ are both 1−correlated. In our arguments we will apply Chernoff-like bounds to negative cylinder dependent variables. We will use the following lemmas.

Lemma 3.4. *Let S_m and S_w be two sets of N_1 and N_2 agents respectively, Suppose that $N_1 \leq N_2$. Let $S_a = \{a_1, a_2, \ldots, a_n\} \subseteq S_m$ and $S_b = \{b_1, b_2, \ldots, b_r\} \subseteq S_w$. Consider a matching between S_m and S_w chosen uniformly at random. Let X_i be an indicator variable which equals 1 if agent a_i is paired with an agent in S_b, and 0 otherwise. Then the set $\{X_i\}$ is negative cylinder dependent and for any $\delta > 0$,*

$$\Pr\left[\sum X_i \geq (1+\delta)\mu\right] \leq e^{-\frac{\delta^2\mu}{3}} \text{ and } \Pr\left[\sum X_i \leq (1-\delta)\mu\right] \leq e^{-\frac{\delta^2\mu}{2}}.$$

Proof. Let $N = \max\{N_1, N_2\}$. Consider any subset $S \subseteq [n]$ where $|S| = k$. W.l.o.g. let $S = [k]$.

Then,

$$E\Big[\prod_{i\in S} X_i\Big] = \Pr\Big[\prod_{i\in S} X_i = 1\Big] = \Pr\left[X_1 = 1, X_2 = 1, \ldots, X_k = 1\right]\Big]$$

$$= \Pr\left[X_1 = 1\right] \cdot \Pr\left[X_2 = 1 | X_1 = 1\right], \ldots, \Pr\left[X_k = 1 | X_1 = 1, X_2 = 1, \ldots, X_{k-1} = 1\right]$$

$$= \frac{r}{N} \cdot \frac{r-1}{N-1}, \ldots, \frac{r-k+1}{N-k+1}.$$

Hence,

$$E\Big[\prod_{i\in S} X_i\Big] \leq \left(\frac{r}{N}\right)^k,$$

while

$$\prod_{i\in S} E[X_i] = \left(\frac{r}{N}\right)^k.$$

Similarly,

$$E\Big[\prod_{i\in S}(1-X_i)\Big] = \Pr\Big[\prod_{i\in S}(1-X_i) = 1\Big] = \Pr\left[X_1 = 0, X_2 = 0, \ldots, X_k = 0\right]\Big]$$

$$= \Pr\left[X_1 = 0\right] \cdot \Pr\left[X_2 = 0 | X_1 = 0\right], \ldots, \Pr\left[X_k = 0 | X_1 = 0, X_2 = 0, \ldots, X_{k-1} = 0\right]$$

$$= \frac{N-r}{N} \cdot \frac{N-r-1}{N-1}, \ldots, \frac{N-r-k+1}{N-k+1}.$$

Hence,

$$E\Big[\prod_{i\in S}(1-X_i)\Big] \leq \left(\frac{N-r}{N}\right)^k,$$

while

$$\prod_{i\in S} E[1-X_i] = \left(\frac{N-r}{N}\right)^k.$$

Thus the set $\{X_i\}$ is negative cylinder dependent. By [Panconesi and Srinivasan 1997, Theorem 3.4] with $\lambda = 1$, Chernoff bounds for sums of independent random variables apply to the sums of negative cylinder dependent random variables as well. This concludes the proof. □

Lemma 3.5. *If $\{X_i\}_{i=1}^n$ are 1-correlated random variables taking value $\{0,1\}$ and $\overline{\mu}$ is an upper bound on $\mu = E[\sum X_i]$, then, for any $\delta > 0$,*

$$\Pr\left[\sum X_i \geq (1+\delta)\overline{\mu}\right] \leq e^{-\frac{\delta^2\overline{\mu}}{3}} \quad \textit{and} \quad \Pr\left[\sum X_i \leq \mu - \delta\overline{\mu}\right] \leq e^{-\frac{\delta^2\overline{\mu}}{2}}.$$

Proof. First we prove that

$$\Pr\left[\sum X_i \geq (1+\delta)\overline{\mu}\right] \leq e^{-\frac{\delta^2\overline{\mu}}{3}}.$$

Let $\overline{\mu} = (1+\theta)\mu$. By Lemma 3.4,

$$\Pr\left[\sum X_i \geq (1+\gamma)\mu\right] \leq e^{-\frac{\gamma^2\mu}{3}}. \tag{3.1}$$

Set $(1+\gamma) = (1+\theta)(1+\delta)$. Then, from equation (3.1) we obtain,

$$\begin{aligned}\Pr\left[\sum X_i \geq (1+\delta)\overline{\mu}\right] &\leq \exp\left[-\frac{(\theta+\delta+\theta\cdot\delta)^2\mu}{3}\right] \leq \exp\left[-\frac{(\delta+\theta\cdot\delta)^2\mu}{3}\right] \\ &\leq \exp\left[-\frac{\delta^2(1+\theta)^2\mu}{3}\right] \leq \exp\left[-\frac{\delta^2(1+\theta)\mu}{3}\right] \leq \exp\left[-\frac{\delta^2\overline{\mu}}{3}\right]\end{aligned}$$

which proves the claim. We will now prove that

$$\Pr\left[\sum X_i \leq \mu - \delta\overline{\mu}\right] \leq e^{-\frac{\delta^2\mu}{2}}$$

Let $\theta = \frac{\mu}{\overline{\mu}}$. By Lemma 3.4,

$$\Pr\left[\sum X_i \leq \mu - \delta\mu\right] \leq e^{-\frac{\delta^2\mu}{2}}$$

Let $\gamma = \frac{\overline{\mu}}{\mu}$. Note that $\gamma \geq 1$

$$\Pr\left[\sum X_i \leq \mu - \delta\overline{\mu}\right] = \Pr\left[\sum X_i \leq \mu - \gamma\delta\mu\right] \leq \exp\left[-\frac{(\gamma\delta)^2\mu}{2}\right]$$
$$\leq \exp\left[-\frac{\delta^2\gamma\overline{\mu}}{2}\right] \leq \exp\left[-\frac{\delta^2\overline{\mu}}{2}\right]$$

which completes the proof.

□

3.5 Lower Bound on the Loss for Any Strategy

The intuition for this result is fairly simple. If an agent remains unmatched for $\sqrt{T}$ steps, then any subsequent proposed match would cause a loss of at least $\sqrt{T}$ to one of the participating agents. Thus to avoid having average losses of $\Omega(\sqrt{T})$, most matches would need to occur during an agent's first $\sqrt{T}$ steps.

But we will show that for at least a constant fraction of the agents, the matches they are offered during their first $\sqrt{T}$ steps will all have the property that the values of the two agents differ by at least $\sqrt{T}$, and consequently one of the participating agents would suffer a $\sqrt{T}$ loss. The overall result follows.

This second claim is not immediate because the probability that an agent is offered a close-in-value match might vary significantly from agent to agent and over time.

We somewhat optimize constants and consequently consider a period of time $w = \Theta(\sqrt{T})$ and value differences w, instead of precisely the value $\sqrt{T}$ used in the above outline.

Proof. (Of Theorem 3.1) We divide the grid into width w columns, where a column includes the low-value side boundary, but not the high-value boundary; one end column may be narrower. We will set the parameter w later.

We consider the set of proposed matches at some arbitrary time t. We say a proposed match is *safe* if the paired agents are in the same or adjacent columns. We also define the male match rate

p_i for column i to be the probability that a man in the column has a safe match. By Lemma 3.3, this is at most the number of women in columns $i-1$, i, and $i+1$ divided by the maximum of the total number of women and the total number of men, which is at most the number of women in these columns divided by the total number of women. Clearly the sum of the male match rates over all the columns is at most 3. The same claim holds for the analogous female match rates.

Consider the men entering the system at time t, which we call the *new* men. Each column contains at most w points at which agents enter the market, namely the points along the column's top edge, and each entering agent is equally likely to be a man or a woman. By applying a Chernoff bound, we see that for any given column i,

$$\Pr\left[\text{\# of new men in column } i \text{ at time } t \geqslant \frac{(1+\delta)nw}{2T}\right] \leqslant e^{-\delta^2 nw/6T}.$$

Applying this bound to every column over τ consecutive time steps yields:

$$\Pr\Bigg[\text{every column receives at most } \frac{(1+\delta)nw}{2T} \text{ new men} \\ \text{for each of } \tau \text{ consecutive time steps}\Bigg] \geqslant \left(1-e^{-\delta^2 nw/6T}\right)^{\tau T/w}.$$

Call this event $\mathcal{E}$. Henceforth we condition on $\mathcal{E}$.

Now suppose that every time an agent was offered a safe match, they accepted it. Recall that p_i is the match rate for column i. By Lemma 3.4, for the new men at time t in column i, for any t,

$$\Pr\left[\text{\# safely matched men } \leqslant \frac{(1+\delta)nwp_i}{2T}\right] \geqslant 1-e^{-\delta^2 nwp_i/6T}.$$

In fact, agents may not accept every proposed safe match; but this only reduces the number of agents safely matched, and therefore the bound on the probability continues to hold.

Furthermore, by Lemma 3.5, letting $\overline{\mu} = \frac{nw\max\{p_i, \frac{w}{T}\}}{2T}$, gives

$$\Pr\left[\text{number of safely matched men } \leqslant \frac{(1+\delta)nw\max\{p_i,\frac{w}{T}\}}{2T}\right] \geqslant 1 - e^{-\frac{\delta^2 nw\max\{p_i,\frac{w}{T}\}}{6T}} \geqslant 1 - e^{-\frac{\delta^2 nw^2}{6T^2}}.$$

Recalling that $\sum_i p_i \leqslant 3$, and applying a union bound over all T/w strips for w successive steps, we obtain, for any given set of new men entering at some time t, over their first w time steps,

$$\Pr\left[\#\text{ of safely matched men} \leq \frac{2(1+\delta)nw^2}{T}\right] \geqslant 1 - \frac{T}{w}we^{-\delta^2 nw^2/6T^2}. \tag{3.2}$$

In addition, for any given set of new men, on applying a Chernoff bound, we know that

$$\Pr\left[\#\text{ of new men} \geq \frac{n(1-\epsilon)}{2}\right] \geqslant 1 - e^{-\epsilon^2 n/4}. \tag{3.3}$$

For each remaining man in each of the first $\tau - w$ sets of new men—of which there are at least $(\tau - w)\left(\frac{n(1-\epsilon)}{2} - \frac{2(1+\delta)nw^2}{T}\right)$— one of the following two cases must apply.

- He has not been matched after spending w time in the system. Now, if and when he is matched, the only way he can avoid suffering a wT loss is to match with a sufficiently higher value woman. In this case the higher value woman suffers at least a wT loss.

- He has been matched within w time but it was not a safe match. In such a match whichever agent had the higher value suffered at least a wT loss.

Since the system runs for τ time steps, this argument can be applied to all agents except those that enter the system during the last w time steps. We deduce that the total loss generated by all these agents is at least $(\tau - w)\left(\frac{n(1-\epsilon)}{2} - \frac{2(1+\delta)nw^2}{T}\right) \cdot wT$.

Note that this loss is being shared by up to $n\tau$ agents. Hence there is an average loss of at least $\frac{1}{2}\left(wT(1-\epsilon) - \frac{4(1+\delta)w^3T}{T}\right) \cdot \frac{\tau-w}{\tau}$. Setting $w = \frac{\sqrt{T}}{4}$, and using the lower bound on τ ($\tau \geq T$), we obtain:

$$\text{average loss per agent} \geq \frac{1}{2}\Big(\frac{T\sqrt{T}(1-\epsilon)}{4} - \frac{T\sqrt{T}(1+\delta)}{16}\Big)\cdot\Big(1-\frac{\sqrt{T}}{4T}\Big).$$

Now we set $\delta = \sqrt{\frac{6T^2}{nw^2}\ln(3n^cT\tau)}$ and $\epsilon = \sqrt{\frac{4}{n}\ln(3\tau n^c)}$. We would like to have $\delta \leq 1$, which we enforce by our choice of constraints on n, T, τ and c (namely $16 \leq T \leq n, c \geq 1, T \leq \tau \leq n^c$ and $n \geq 96T(2c+2)\ln n$). These constraints also ensure that $\epsilon \leq 1/16$. Substituting $\delta \leqslant 1$ and $\epsilon \leq 1/16$ yields:

$$\text{average loss per agent} \geq \frac{7T\sqrt{T}}{128}\cdot\frac{15}{16} \geq \frac{T\sqrt{T}}{20}.$$

By (3.2) and (3.3), this bound holds with probability at least

$$\begin{aligned}
\Pr[\mathcal{E}]\cdot\Big(1-\tau Te^{-\frac{\delta^2nw^2}{6T^2}}-\tau e^{-\frac{\epsilon^2 n}{4}}\Big) &\geqslant \Big(1-e^{-\frac{\delta^2nw}{6T}}\Big)^{\tau\frac{T}{w}}\cdot\Big(1-\tau Te^{-\frac{\delta^2nw^2}{6T^2}}-\tau e^{-\frac{\epsilon^2 n}{4}}\Big)\\
&\geq \left[1-\Big(\frac{1}{3n^cT\tau}\Big)^{4\sqrt{T}}\right]^{4\tau\sqrt{T}}\cdot\Big(1-\frac{1}{3n^c}-\frac{1}{3n^c}\Big)\\
&\geq \left[1-4\tau\sqrt{T}\Big(\frac{1}{3n^cT\tau}\Big)^2\right]\cdot\Big(1-\frac{1}{3n^c}-\frac{1}{3n^c}\Big)\\
&\geq \left[1-\Big(\frac{1}{3n^c}\Big)^{4\sqrt{T}}\right]\cdot\Big(1-\frac{2}{3n^c}\Big) \geq \Big(1-\frac{1}{n^c}\Big).
\end{aligned}$$

□

3.6 Upper Bound on the Loss when Using the Modified Reasonable Strategy

The lower bound suggests that plausible agent strategies will yield a constant probability of matching every $\sqrt{T}$ steps. This would imply that the number of agents present decreases geometrically with agent age; more precisely, there would be a constant factor decrease for every $\sqrt{T}$ increment in age. Then, in order to maintain match probabilities, all agents would have to be willing to match with young agents who will accept them. In fact, the decreases we just described are far from uniform, which makes the analysis quite non-trivial. Nonetheless, the above intuition informed the design of the following agent strategies. The first strategy, which we call "a reasonably good strategy" seems quite natural, but for ease of analysis we consider a modified strategy which we prove to be asymptotically within a constant factor of optimal.

We define the *worth* of an agent to be $v_i \cdot (T - t_i)$; this is the maximum utility its partner could derive from a match with this agent. Note that the worth of an agent decreases as it ages.

A Reasonably Good Strategy In this strategy an agent accepts a proposed match if it gives the agent utility at least $v_i \cdot (T - t_i) \cdot (1 - \frac{1}{\sqrt{T}} - \frac{t_i}{T})$. The terms $1/\sqrt{T}$ and t_i/T are present to approximately balance the expected loss of utility from not matching in a single step with the marginal gain in utility agent i could receive from being more demanding in terms of the minimum worth it will accept in a partner.

The Modified Reasonable Strategy We partition the $T \times T$ size space into the regions defined below, as shown in Figure 3.1. In the modified strategy, an agent accepts a proposed match exactly if the proposed partner lies in the same region. This partition uses regions of two kinds, which we call *strips*.

- Type 1 strips: these are strips that have new people entering the strip at the top. The i-th Type 1 strip is defined as the region between the parallel lines $v = 2(t-1) + T + (i-1)\sqrt{T}$ and $v = 2(t-1) + T + i\sqrt{T}$; they have $\sqrt{T}$ width and $\sqrt{T}/2$ height. Points on the first (left) line are included in the strip, but points on the second (right) line are excluded. There are $\sqrt{T}$ Type 1 strips.

- *Type* 2 *strips*: these strips do not touch the top boundary of the box. The strips are again defined by parallel lines. They have successive heights $\sqrt{T}$, $\sqrt{T}$, $2\sqrt{T}$, and then repeatedly doubling up to $T/2$. Here the points on the first (upper) line are excluded from the strip and the points on the second (lower) line are included in the strip. There are $\log_2 \sqrt{T} + 1$ Type

2 strips.

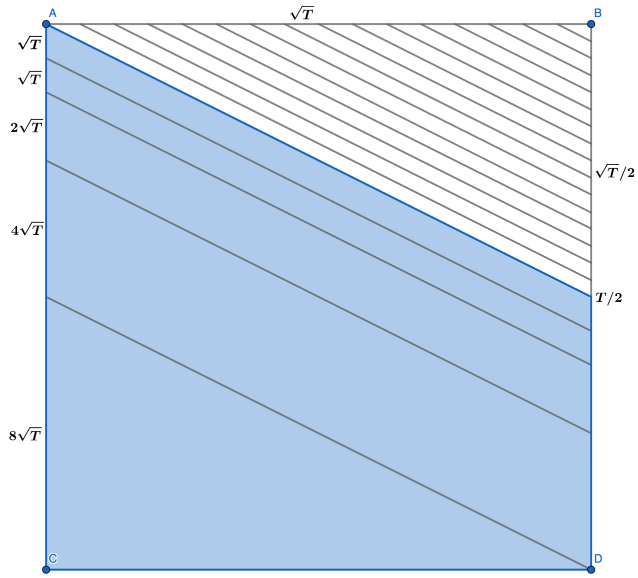

Figure 3.1: The two types of strips used to partition the matching pool.

We note that with the previously stated reasonable strategy, agents would be willing to match with some agents outside their strip and would reject some agents in the same strip. However, using the modified strategy simplifies the analysis, for if all agents use the modified strategy, agents will definitely get accepted when they accept a match. We will prove that the modified strategy is not much worse than the optimal strategy in terms of the average loss of value suffered by an agent.

Outline of the proof of the upper bound Our analysis assumes the following constraints on n and T.

$$\begin{aligned} &c \geq 1, \\ &T \geqslant 676, \\ &n \geq (3654 + 2436e^{12} + 546(e^{12}+1)c)^2(3c+4)T^3(\log_2 n)^2 \ln n. \end{aligned} \tag{3.4}$$

The result follows from a high-probability inductive bound on the overall population, the strip populations, and the male-female imbalances in each strip. We start at time $t = 0$. Time t will refer to the moment after the new agents have entered in this step, but before the match occurs.

Lemma 3.6. *Let N denote the total number of strips. Suppose that the constraints in (3.4) hold. Then, with probability at least $1 - 1/n^c$, the following inductive hypothesis $H(t)$ holds at the start of every time step t, immediately following the entry of the new agents at time t, for $\sqrt{T} \leqslant t \leqslant n^c$.*

1. *The total population is at most $\frac{3}{2}nN + n$.*

2. *The population of every Type 1 strip is at most $2.6n$.*

3. *The population of every Type 2 strip is at most $\frac{7.5n\sqrt{T}}{\text{maximum height of the strip}}$.*

4. *The population in the bottommost Type 2 strip is no more than $60n/\sqrt{T}$.*

5. *In every strip s, except possibly the bottommost Type 2 strip, the imbalance, $\mathrm{Imb}(s,t) = \left|\text{the number of men in } s - \text{the number of women in } s\right| \leq n/25\sqrt{T}$.*

Proof. (Sketch.) We will show in Theorems 3.10 and 3.12–3.15 that each of the above five clauses holds with high probability. The last of these results also requires a high-probability lower bound, Theorem 3.9, on the population size in the same time range. In addition, in Theorem 3.16, we show that, with high probability, the inductive hypothesis is true initially. Summing the failure probabilities prove the lemma. This calculation can be found in Appendix A.1.6. □

With this result in hand we can upper bound the average agent loss.

3.6.1 The Theorems and Proof Sketches

Let $\widetilde{\mathcal{E}}$ be the event that the inductive hypothesis $H(t)$ holds at the start of time step t immediately following the arrival of the new agents in this step, for $\sqrt{T} \leqslant t \leqslant n^c$.

3.6.1.1 Bounding the loss

We first bound an individual agent's loss based on its match time. We then obtain an overall bound on the loss. As argued below, Theorem 3.2 follows immediately.

Lemma 3.7. *In the modified reasonable strategy, if an agent with value v matches at time t, its utility loss is at most $4Tt + 2t\sqrt{T}$.*

This result follows by a simple calculation based on the strip geometry. The proof is in Appendix A.1.1.

Theorem 3.8. *Suppose the constraints in (3.4) hold. Also, suppose that all agents follow the modified reasonable strategy. In addition, suppose the system runs for $\tau \geqslant 2T$ time steps, where $\tau \leq n^c$. Then the average loss per departing agent over these τ steps will be at most $11T\sqrt{T}$.*

Proof. Consider the first τ time steps of the matching process. Let n_i denote the number of agents who match and thereby leave the pool at age i during these τ steps. By Lemma 3.7, each such agent suffers a loss of at most $4Ti + 2T\sqrt{T}$. Thus the total loss is bounded by:

$$\text{Total loss} \leqslant \sum_{i=0}^{T-1} \left(4Ti \cdot n_i + 2T\sqrt{T} \cdot n_i\right).$$

Each agent who is matched at age i is present in the matching pool for $i+1$ steps. By clause 1 of the inductive hypothesis in Lemma 3.6, at each time during this period, the population of the matching pool is at most $\frac{3}{2}nN + n \leq \frac{3}{2}n(\sqrt{T} + \log_2 \sqrt{T} + 1) + n \leq 2n\sqrt{T}$, where the last inequality follows from $\sqrt{T} \leqslant 26$ due to constraint (3.4). Thus,

$$\sum_{i=0}^{\tau-1} (i+1)n_i \leqslant 2n\sqrt{T} \cdot \tau.$$

Therefore,

$$\text{Total loss} \leqslant 8nT\sqrt{T} \cdot \tau + \sum_{i=0}^{T-1} 2T(\sqrt{T} - 2) \cdot n_i.$$

Let $D \triangleq \sum_{i=0}^{\tau-1} n_i$, the number of agents that leave during the first τ steps. We observe that D is at most $n\tau$, the number of agents that entered during this period. Also, as the population of the pool at any time is at most $2n\sqrt{T}$, we see that $D \geq n\tau - 2n\sqrt{T}$. By assumption, $\tau \geqslant 2T$ and $\sqrt{T} \geqslant 26$, so

$$\frac{12}{13} n\tau \leqslant D \leqslant n\tau.$$

This yields the following bound on the total loss:

$$\text{Total loss} \leq 8n\tau T\sqrt{T} + 2n\tau T\sqrt{T} \leqslant 10n\tau T\sqrt{T}.$$

And therefore,

$$\text{Average loss per agent} = \frac{\text{Total loss}}{D} \leq \frac{10nT\tau\sqrt{T}}{\frac{12}{13}n\tau} < 11T\sqrt{T}.$$

□

Proof. (Of Theorem 3.2) This follows immediately from Lemma 3.6 and Theorem 3.8. □

23

3.6.1.2 Total Size Lower Bound

Theorem 3.9. *Suppose $H(t)$ and the constraints in (3.4) hold. If all agents follow the modified reasonable strategy, then with probability at least $1 - 1/n^{2c+1}$, for every time $t \in [\sqrt{T}, n^c]$, the population in the matching pool is at least $\frac{1}{3}n\sqrt{T}$.*

Proof. (Idea.) We consider only the new agents that entered the matching pool over the last $\sqrt{T}$ time steps. We then bound how many of these agents could have been matched in this time period. Suppose that at any particular time step t, the match rate experienced by the men in strip i is p_i. The critical observation is that the sum of the p_i is at most 1. The same is true for the women. This allows us to prove that even if we could set the match rates in an adversarial manner, only about $n/\sqrt{T}$ of the agents that entered at any one time could be matched in any single time step (in the discussion here, we neglect the effects of variance). This allows us to show that, of the agents we consider, only about $\sum_{i=1}^{\sqrt{T}} in/\sqrt{T} \approx n\sqrt{T}/2$ could have been matched over the last $\sqrt{T}$ time steps. This provides a lower bound on the total size of roughly $n\sqrt{T}/2$. Accounting for the variance that can occur when achieving a high probability bound causes the bound on the number of matches to degrade to $n\sqrt{T}/3$. The full proof can be found in Appendix A.1.2. □

3.6.1.3 Population Upper Bound

Theorem 3.10. *Suppose $H(t)$ and the constraints in (3.4) hold. If all agents follow the modified reasonable strategy, then at the start of time step $t+1$, with probability at least $1-1/n^{2c+1}$, the total population of the matching pool will be at most $(3/2)nN+n$, where N is the total number of strips.*

Proof. (Idea.) We seek to lower bound the number of matches in one time step. If it exceeds the number of incoming agents, then the total population reduces. The expected number of matches is minimized when the strip populations are equal, and on applying Lemma 3.3, this yields the following lower bound on the number of matched women (or men): $[N \cdot (P/2N)^2/(P/2)] = P/(2N)$, where P is the upper bound on the population. This yields the condition $P/N \leqslant n$, or $P \leqslant nN$. The argument is completed by taking account of the deviations needed to ensure a high-probability bound. The full proof can be found in Appendix A.1.3. □

3.6.1.4 Upper Bound on the Size of a Strip.

We begin with a technical lemma.

Lemma 3.11. *Let s be a strip, and let S be an arbitrary subset of the men and women in s. Let m be the number of men and w be the number of women in S. In addition, let X be the imbalance for the whole of s. Then the expected number of people in S that are matched in a single step is at least*

$$\frac{\frac{(m+w)^2}{2} - \frac{X^2}{2}}{\max\{\#\ \textit{of men, \# of women}\}\ \textit{in the whole population}}$$

Proof. We need only consider the case that $|X| \leq m + w$. [2] Let m_t denote the total number of men in this strip and w_t the total number of women. In addition, let $\Delta \triangleq m - w$, $P \triangleq m + w$, $Q \triangleq m_t + w_t$ and $X = m_t - w_t$. Then, $m = \frac{P+\Delta}{2}$, $w = \frac{P-\Delta}{2}$, $m_t = \frac{Q+X}{2}$ and $w_t = \frac{Q-X}{2}$. The expected number of people matched in this subset of men and women is

$$\frac{mw_t + m_t w}{\max\{\#\text{ of men, \# of women}\}\text{ in the whole population}}.$$

We now focus on the numerator: $mw_t + m_t w = (PQ - X\Delta)/2 = (P^2 + P(Q-P) - X^2 - X(\Delta - X))/2$. In order to show this is larger than $\frac{P^2}{2} - \frac{X^2}{2}$, it suffices to show $P(Q-P) \geq X(\Delta - X)$.

As $m_t \geq m$ and $w_t \geq w$, $Q - P \geq \Delta - X$ and $Q - P \geq X - \Delta$. Recall that it suffices to consider the case $|X| \leq m + w = P$. Combining these two inequalities yields $P(Q-P) \geq X(\Delta - X)$, which proves the result. □

Next, we give an upper bound on the size of a Type 1 strip.

Theorem 3.12. *Suppose $H(t)$ and the constraints in (3.4) hold. If all agents follow the modified reasonable strategy, then at time $t+1$, right after the new agents have entered, with probability $> 1 - \frac{1}{n^{2c+1}}$, each Type 1 strip will continue to have population at most dn, where $d = 2.6$.*

[2] Otherwise, the bound is negative.

Proof. (Sketch). Consider a strip s' and its successor strip s (the strip immediately to its left). We will follow the collection of agents occupying $\sqrt{T}$ adjacent diagonals over $\sqrt{T}$ steps, beginning with the at most dn agents in strip s' and ending in strip s, with the remainder of these agents plus any new agents who have entered these diagonals. The heart of our proof is to show that in a single step we maintain the dn bound on the number of agents in this collection of advancing diagonals. The basic idea is straightforward: we compute a lower bound on the expected number of matches using Lemma 3.11 taking into account the maximum possibly imbalance, add the incoming agents and correct for variance. One more important detail is that the expected number of matches is minimized if, in the collection of agents we are tracking, half are in strip s and half are in s'; so this is the value we use in these calculations. The full proof can be found in Appendix A.1.4. □

Theorem 3.13. *Suppose $H(t)$ and the constraints in (3.4) hold. If all agents follow the modified reasonable strategy, then at time $t+1$, right after the new agents have entered, with probability $> 1 - \frac{1}{n^{2c+1}}$, each Type 2 strip (apart from the bottommost one) will continue to have population at most $\frac{gn\sqrt{T}}{\textit{height of the strip}}$ where $g = 7.5$.*

Proof. (Idea.) For the topmost Type 2 strip s we obtain a bound of $2 \cdot 2.6n = 5.2n\sqrt{T}/\sqrt{T}$, as the items in s are obtained from its predecessor strip over the previous $\sqrt{T}$ steps, i.e. the sum of the contents at times $\sqrt{T}/2$ and $\sqrt{T}$ earlier. The same bound applies to the second Type 2 strip. Each subsequent Type 2 strip s has twice the height of its predecessor. Let s have height H. The contents of s come from its predecessor over a period of length H, which by the inductive hypothesis contain at most $2gn\sqrt{T}/(H/2) = 4gn\sqrt{T}/H$ agents. To prove our bound, we need to show at least $3gn\sqrt{T}/H$ of them are removed during these H steps. Again, as in Theorem 3.12, we seek to track a population as it moves from s' to s. The challenge is that in the analysis this population shrinks over time and the match rate is proportional to the square of this population. To get a fairly tight bound, we formulate this as a differential expression and determine the smallest value for the constant g that enables this number of matches. The full proof can be found in Appendix A.1.5. □

Theorem 3.14. *Suppose $H(t)$ and the constraints in (3.4) hold. If all agents follow the modified reasonable strategy, then at time $t+1$, right after the new agents have entered, the strip population for the bottommost Type 2 strip will continue to be at most $\frac{60n}{\sqrt{T}}$.*

This simple calculation is deferred to Appendix A.1.5.

3.6.1.5 Bound on Imbalance

Theorem 3.15. *Suppose that $H(\tau)$ and the constraints in (3.4) hold. If all agents follow the modified reasonable strategy, then with probability at least $1 - 2/n^{2c+1}$, in every strip s (except possibly the bottommost Type 2 strip),* $\mathrm{Imb}(s) \leqslant n/25\sqrt{T}$.

Proof. We divide each strip into thin diagonals of width 1. Let the diagonal include the bottom but not the top boundary. Notice that for each value, a diagonal contains at most one grid point.

We introduce the following notation w.r.t. diagonal d at time step τ, where we are conditioning on the outcome of step $\tau - 1$.

$$
\begin{aligned}
I(d,\tau) &= E[(\text{number of men at time } \tau - \text{number of women at time } \tau)] \\
X(d,\tau) &= (\text{number of men matching at time } \tau - \text{number of women matching at time } \tau) \\
&\quad - E[(\text{number of men matching at time } \tau - \text{number of women matching at time } \tau)] \\
Y(d,\tau) &= \text{number of men entering at time } \tau - \text{number of women entering at time } \tau \\
A(d,\tau) &= (\text{number of men matching at time } \tau + \text{number of women matching at time } \tau)/2.
\end{aligned}
$$

$I(d,\tau)$ is measured after the entry of the new agents at time τ but prior to the match for this step. Also, note that $Y(d,\tau) = 0$ if d is in a Type 2 strip.

In addition, observe that the imbalance $\mathrm{Imb}(s)$ at the start of step t equals $\sum_{d\in s} I(d,t)$.

We observe that a match between two agents in distinct diagonals of the same strip will increment the (number of men $-$ number of women) in one diagonal and decrement it in the other. Thus there is a zero net change over all the diagonals in the strip due to the matches. However, as the agents all age by 1 unit during a step, some agents enter the strip and some leave, which can cause changes to the imbalance within a strip. However, the entry of new agents can introduce new imbalances. We will need to understand more precisely how these imbalances evolve.

It is convenient to number the diagonals as $d_1, d_2, d_3, \ldots$, in right to left order.

Claim 3.6.1. *Let d_i and d_j be two diagonals in the same strip s. For brevity, let $I_i \triangleq I(d_i, \tau-1)$, $I_j \triangleq I(d_j, \tau-1)$, $A_i \triangleq A(d_i, \tau-1)$, $A_j \triangleq A(d_j, \tau-1)$, $X_i \triangleq X(d_i, \tau-1)$, $X_j \triangleq X(d_j, \tau-1)$. Finally, let R denote the maximum of the total number of men and the total number of women in the system at time $\tau-1$.*

Then the new imbalance on diagonal d_i, prior to every unmatched agent adding 1 to their age (which causes the agents on d_i to move to d_{i+1}), denoted by $I'(d_i, \tau)$, is given by:

$$I'(d_i, \tau) = \\ I_i + X_i - \sum_{d_j \in s} \left[X_i \frac{(2A_j - I_j - X_j)}{4R} - X_j \frac{(2A_i - I_i - X_i)}{4R} + I_i \frac{(2A_j - I_j - X_j)}{4R} - I_j \frac{(2A_i - I_i - X_i)}{4R} \right];$$

and $I(d_i, \tau) = I'(d_{i-1}, \tau-1) + Y(d, \tau)$.

This claim is shown by considering the expected number of matches involving agents in diagonals d_i and d_j. The proof can be found in Appendix A.1.6.

The expression $X_i(2A_j - I_j - X_j)/4R$ reflects the reduction of the contribution of X_i to the total imbalance on diagonal d_i and the corresponding increase on diagonal d_j. Thus it is convenient to view the multiplier $(2A_j - I_j - X_j)/4R$ as indicating the fraction of X_i that is being moved to diagonal j; the remaining fraction of X_i remains on d_i.

$X(d,\tau)$ and $Y(d,\tau)$ are generated at diagonal d at time τ. In each subsequent time step the portion on each diagonal where it is present will be further redistributed:

1. Due to the expected matching at time $\tau' \geqslant \tau$, each portion of $X(d,\tau)$ and $Y(d,\tau)$ spreads to other diagonals in the same strip.
2. At the end of time step τ' the portions of $X(d,\tau)$ and $Y(d,\tau)$ present on diagonal d_i move to diagonal d_{i+1}.

Building on these observations, we will show our bound on the imbalance by means of the following two arguments. Specifically, we show that:

1. For any τ and τ', the total contribution from $X(\cdot,\tau)$ and $Y(\cdot,\tau)$ to strip s at time τ' is bounded.
2. For times $\tau' \geqslant \tau + \Omega(T \log n)$, the remaining portions of $X(\cdot,\tau)$ and $Y(\cdot,\tau)$ in the market are small.

Bound on the contribution of X to the strip s Notice that $\sum_{d_i \in s} I'(d_i,\tau) = \sum_{d_i \in s} I(d_i,\tau-1)$, for the coefficients multiplying X_i cancel, as they also do for I_i. Thus we can think of this process as redistributing the imbalance, but not changing the total imbalance.

Over time an imbalance $X(d_i,\tau)$ will be redistributed over many diagonals. We write $X(d_i,\tau,d_j,\tau')$ to denote the portion of $X(d_i,\tau)$ on diagonal d_j at time τ'. d_j need not be in the same strip as d_i. Note that $\sum_{d_j} X(d_i,\tau,d_j,\tau') = X(d_i,\tau)$ for all $\tau' \geqslant \tau$. $Y(d_i,\tau,d_j,\tau')$ is defined analogously.

An important property concerns the relative distribution of the $X(d_i,\tau,d_j,\tau')$ and the $X(d_k,\tau,d_j,\tau')$. In a sense made precise in the following claim, if $k > i$ the d_k terms remain to the left of the d_i terms.

For the purposes of the following claim, we treat the final strip as a single diagonal, and in addition ignore the fact that people depart at age T (which means that once an imbalance appears in this strip it remains there). The reason this strip is different is that it covers the whole of the bottom boundary and so is the only strip from which people leave the system by aging out.

Claim 3.6.2. *For all ℓ, for all $i < k$, and for all $\tau' \geqslant \tau$, $\left|\sum_{j>\ell} X(d_i, \tau, d_j, \tau')\right| \leqslant \left|\sum_{j>\ell} X(d_k, \tau, d_j, \tau')\right|$. The same property holds for the $Y(d_i, \tau, d_j, \tau')$.*

Proof. We prove the result for the X terms by induction on τ'; the same argument applies to the Y terms. Clearly the property holds for $\tau' = \tau$. Let $x_{ij} \triangleq X(d_i, \tau, d_j, \tau')/X(d_i, \tau)$, and define x_{kj} analogously. Our claim states that $\sum_{j>\ell} x_{ij} \leqslant \sum_{j>\ell} x_{kj}$; we need to show it holds at time $\tau'+1$ also. We view the x_{ij} as sitting on the unit interval, with x_{ij} taking a portion of length x_{ij}, ordered by increasing j, and likewise for the x_{kj}. We map aligned portions of the x_{ij} and $x_{kj'}$ to each other. This mapping has the property that the j index in the x_{ij} term is always equal to or smaller than the j' index in the $x_{kj'}$ term.

Let's look at how aligned portions of x_{ij} and $x_{kj'}$ are dispersed in the next step. If they are in distinct strips, then $j < j'$ and this property is maintained for all the dispersed portions.

We view the multiplier $(2A_j - I_j - X_j)/4R$ in Claim 3.6.1 as specifying the fraction of X_i that moves from diagonal i to diagonal j. Notice that this multiplier is the same for every diagonal in this strip.

We also note that I_i consists of a sum of terms $X(d_i, \tau, d_j, \tau')$ and $X(d_i, \tau, d_j, \tau')$ for diagonals d_j in the same strip as d_i or to the right of d_i. Furthermore, the multiplier $(2A_j - I_j - X_j)/4R$ specifies the fraction of each of these terms that moves from diagonal i to diagonal j. Thus if d_j and $d_{j'}$ are in the same strip, the X terms corresponding to the aligned portions of x_{ij} and $x_{kj'}$ are redistributed identically, thereby maintaining the property for these fragments. Naturally, the property also continues to hold for undispersed fragments.

Finally, shifting down by one diagonal, as is done following the dispersal, will leave the property unaffected.

□

Later, we will show a common bound B on the sums $\left|\sum_{i\leqslant j\leqslant k} X(d_j,\tau)\right|$, which holds for all d_i and d_k in the same strip and all[3] τ. With this bound and Claim 3.6.2 in hand, for each strip s, we can bound the contribution of the $X(d_i,\tau,d_j,\tau')$ summed over all d_i and over $d_j \in s$ by $2B$.

Claim 3.6.3. *For all $\tau' \geqslant \tau$, for every strip s, $\left|\sum_{d_i;d_j\in s} X(d_i,\tau,d_j,\tau')\right| \leqslant 2B$.*

Proof. Let $d_{r(s)}$ be the rightmost

(lowest index) diagonal in s and $d_{l(s)}$ be the leftmost (highest index) diagonal in s. Let $w_i = \sum_{j\geqslant r(s)} X(d_i,\tau,d_j,\tau')/X(d_i,\tau)$. Let's consider $\sum_{d_i\in s';j\geqslant r(s)} X(d_i,\tau,d_j,\tau') = \sum_{d_i\in s'} w_i \cdot X(d_i,\tau)$. Notice that $\sum_{r(s')\leq i\leq l(s)} X(d_i,\tau) = 0$. By Claim 3.6.2, $w_i \leqslant w_k$, for $i < k$. Thus,

$$\left|\sum_{d_i\in s';j\geqslant r(s)} X(d_i,\tau,d_j,\tau')\right| = \left|\sum_{d_i\in s'} w_i \cdot X(d_i,\tau)\right| \leq \sum_{r(s')\leqslant r<l(s')} (w_i - w_{i-1})\left|\sum_{r\leqslant i\leqslant l(s')} X(d_i,\tau)\right|$$
$$\leqslant (w_{l(s')} - w_{r(s')}) \cdot \max_{r\geq r(s')}\left|\sum_{r\leqslant i\leqslant l(s')} X(d_i,\tau)\right|.$$

We apply this bound to the diagonals from every strip to obtain:

$$\left|\sum_{d_i;j\geqslant r(s)} X(d_i,\tau,d_j,\tau')\right| = \left|\sum_{s'}\sum_{d_i\in s';j\geqslant r(s)} X(d_i,\tau,d_j,\tau')]\right| \leq \sum_{s'}(w_{l(s')} - w_{r(s')}) \cdot B \leq B. \tag{3.5}$$

Using the same argument, $\left|\sum_{d_i;j\geqslant l(s)+1} X(d_i,\tau,d_j,\tau')\right| \leqslant B$, since $l(s)+1 = r(s'')$ where s'' is the strip immediately below s. Therefore,

$$\left|\sum_{d_i;d_j\in s} X(d_i,\tau,d_j,\tau')\right| = \left|\sum_{d_i;j\geqslant r(s)} X(d_i,\tau,d_j,\tau') - \sum_{d_i;j\geqslant l(s)+1} X(d_i,\tau,d_j,\tau')\right| \leqslant 2B.$$

□

Claim 3.6.4. *For any time $\tau \leq n^c$, with probability at least $1-\frac{1}{n^{2c+1}}$, $B \leq 96\left[\frac{n\ln\left(4n^{3c+1}(T^2/32+T/8)N\right)}{\sqrt{T}}\right]^{1/2}$.*

[3]The calculation for the bound proved in Claim 3.6.4 only applies to $|\sum_{i\leqslant j\leqslant k} X(d_j,\tau)|$, where $\tau > \sqrt{T}$. However for times in the initial $\sqrt{T}$ steps, the bound is only better. A calculation of this bound for times in this initial period is done in the proof of Theorem 3.16; see Claim A.1.5 in Appendix A.1.7.

Proof. First we bound $|\sum_{d\in S} X(d,\tau)|$ for any subset S of consecutive diagonals in a strip s. Suppose the total number of men in S is m and the total number of women is w.

By Theorem 3.9, the total population is at least $1/3\cdot n\sqrt{T}$. By Theorem 3.10, it is at most $3nN/2+n$. In addition, by the inductive hypothesis, the total imbalance is bounded by the bottommost strip population plus the individual strip imbalances, and this is at most $60n/\sqrt{T} + 25nN/\sqrt{T}$. Therefore,

$$\frac{n\sqrt{T}}{6} \le \max\left\{\begin{array}{c}\text{total number of men,}\\ \text{total number of women}\end{array}\right\} \le \frac{1}{2}\left(\frac{3n(\sqrt{T}+\log_2\sqrt{T}+1)}{2}+n+60n/\sqrt{T}+nN/25\sqrt{T}\right).$$

As $\sqrt{T} \ge 26$ by constraint (3.4),

$$\frac{n\sqrt{T}}{6} \le \max\left\{\begin{array}{c}\text{total number of men,}\\ \text{total number of women}\end{array}\right\} \le n\sqrt{T}. \tag{3.6}$$

Let $M = \max\{\text{total number of men, total number of women}\}$. Lemmas 3.4 and 3.3 yield the following bound on the deviation from the expected number of the number of men in S matched in a given time step:

$$\Pr\left[\left|\begin{array}{c}\text{number of men matched}\\ -E[\text{number of men matched}]\end{array}\right| > \frac{mw\epsilon}{M}\right] \le 2e^{-mw\epsilon^2/3M}. \tag{3.7}$$

By the lower bound on M provided by (3.6):

$$\begin{aligned}&\Pr\left[\left|\begin{array}{c}\text{number of men matched}\\ -E[\text{number of men matched}]\end{array}\right| > \frac{6mw\epsilon}{n\sqrt{T}}\right]\\ &\le \Pr\left[\left|\begin{array}{c}\text{number of men matched}\\ -E[\text{number of men matched}]\end{array}\right| > \frac{mw\epsilon}{M}\right].\end{aligned}$$

And by the upper bound on M given by (3.6), $2e^{-mw\epsilon^2/3M} \leq 2e^{-mw\epsilon^2/3n\sqrt{T}}$.

We now apply these two bounds to equation (3.7) to obtain:

$$\Pr\left[\left| \begin{array}{c} \text{number of men matched} \\ -E[\text{number of men matched}] \end{array} \right| > \frac{6mw\epsilon}{n\sqrt{T}}\right] \leq 2e^{-2mw\epsilon^2/9n\sqrt{T}}.$$

The same reasoning can be applied to the number of women matched in S.

We set $\epsilon = \left[\frac{3n\sqrt{T}}{mw}\ln\left(4n^{3c+1}(T^2/32 + T/8)N\right)\right]^{1/2}$. By the inductive hypothesis, $m + w \leq 7.5n$, and therefore $mw \leq (15n/4)^2$. We obtain:

$$\frac{6mw\epsilon}{n\sqrt{T}}\left[\frac{3mw\ln\left(4n^{3c+1}(T^2/32 + T/8)N\right)}{n\sqrt{T}}\right]^{1/2} = \frac{45\sqrt{3}}{2}\left[\frac{n\ln\left(4n^{3c+1}(T^2/32 + T/8)N\right)}{\sqrt{T}}\right]^{1/2},$$
$$\text{and} \quad 2e^{-2mw\epsilon^2/9n\sqrt{T}} \leqslant \frac{1}{2n^{3c+1}(T^2/32 + T/8)N}.$$

On adding the bounds for the numbers of men and women, this yields:

$$\Pr\left[\left|\sum_{d\in S} X(d,\tau)\right| \leq 45\sqrt{3}\left[\frac{n\ln\left(4n^{3c+1}(T^2/32 + T/8)N\right)}{\sqrt{T}}\right]^{1/2}\right] \leq \frac{1}{n^{3c+1}(T^2/32 + T/8)N}. \tag{3.8}$$

Recall that there are N strips, at most n^c rounds, and, for each strip, there are at most $(T^2/32 + T/8)$ choices of l and r. Therefore, the total failure probability is at most $\frac{1}{n^{2c+1}}$.

□

Bound on the contribution of Y to strip s. As for X, we define $Y(d_i, \tau, d_j, \tau')$ to be the portion of $Y(d_i, \tau)$ on diagonal d_j at time τ'.

Claim 3.6.5. *With probability at least* $1 - \frac{1}{n^{2c+1}}$, *for all* $\tau' \geqslant \tau$, *for every strip* s, $\left|\sum_{d_i;d_j\in s} Y(d_i,\tau,d_j,\tau')\right| \leqslant 2\sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)}$.

The proof of this claim is similar in spirit to that of Claim 3.6.4. We defer it to Appendix A.1.6.

Remaining X and Y in the market. Next, we want to show that after $O(T)$ time the portions of X and Y remaining in the market are small.

Claim 3.6.6. *$\frac{e^2 \ln 2}{\log_2(4/3)}\sqrt{T}(\sqrt{T}+\log_2(2n^k))$ time after their creation, there is only a $\frac{1}{2n^k}$ fraction of $X(d,\tau)$ and $Y(d,\tau)$ remaining in the Type 1 strips.*

Proof. Consider some $X(d,\tau)$ or $Y(d,\tau)$ generated in a Type 1 strip.

We first bound $\sum_{j:d_j\in s}(2A_j - I_j - X_j)/4R$ for any Type 1 strip s. By Theorem 3.9, the total size of the population is lower bounded by $(1/3)n\sqrt{T}$. By the inductive hypothesis, any Type 1 strip s has total size at most $2.6n$. The term $\sum_{j:d_j\in s}(2A_j - I_j - X_j)$ is 2 times the total number of women in strip s. By the inductive hypothesis, the number of women in s is at most $1.3n + n/50\sqrt{T}$. Lemma 3.3 provides the following upper bound on the probability that a man receives a match in a Type 1 strip:

$$\sum_{j:d_j\in s}\frac{(2A_j - I_j - X_j)}{4R} \leq \frac{1}{2}\cdot\frac{1.3n+\frac{n}{50\sqrt{T}}}{\frac{1}{6}n\sqrt{T}} < \frac{4}{\sqrt{T}}, \quad \text{as } (\sqrt{T}\geqslant 26 \text{ by constraint 3.4}). \tag{3.9}$$

Consider any $X(d,\tau,d',\tau')$. If d' is in a Type 1 strip then by (3.9) in one step at most $\frac{4}{\sqrt{T}}$ of it disperses to some location in the same strip, and at least $1-\frac{4}{\sqrt{T}}$ of it moves down distance one. This implies that in $\sqrt{T}/2$ time a Type 1 strip loses at least e^{-2} of the $X(d,\tau,d',\tau')$ that had been present within it at time τ'. Let $K_1 = e^2\ln 2$. By time $\tau' + K_1\sqrt{T}/2$ at least half of the $X(d,\tau,d',\tau')$ in a Type 1 strip has moved out of the strip.

We number the Type 1 strips from top to bottom. Let γ be the distribution of $X(d,\tau)$ (or $Y(d,\tau)$) where γ_i is the fraction of $X(d,\tau)$ (or $Y(d,\tau)$) in strip i. Recall that there are $\sqrt{T}$ Type 1 strips. We consider the worst case: the $X(d,\tau)$ starts out in the topmost strip. Define a potential function $\phi(\gamma) = \sum_{i=1}^{\sqrt{T}}\gamma_i\cdot 2^{\sqrt{T}-i+1}$. Any fraction of $X(d,\tau)$ that has left the bottommost Type 1 strip contributes nothing to the potential. The initial potential is $2^{\sqrt{T}}$. Every $K_1\sqrt{T}$ time steps, the potential decreases by at least $1/4$. Therefore, after $\frac{1}{\log_2(4/3)}K_1\sqrt{T}\log_2(2^{\sqrt{T}}2n^k)$ time, the potential would have reduced to at most $\frac{1}{2n^k}$, which means that the fraction of $X(d,\tau)$ (or $Y(d,\tau)$) in the Type 1 strips after $\frac{1}{\log_2(4/3)}K_1\sqrt{T}(\sqrt{T}+\log_2(2n^k))$ time is at most $\frac{1}{2n^k}$. □

We will analyze the progress through the Type 2 strips, apart from the bottommost one, in a similar way. The proof can be found in Appendix A.1.6.

Claim 3.6.7. $\frac{e^2 \ln 2}{\log_2(4/3)}\sqrt{T}(\sqrt{T}+\log_2(2n^k))+\frac{e^{12}\ln 2}{4\log_2(4/3)}T\log_2(2n^k\sqrt{T})$ *time after their creation, there is only* $\frac{1}{n^k}$ *fraction of* $X(d,\tau)$ *and* $Y(d,\tau)$ *remaining in any strip other than the bottommost Type 2 strip.*

The Total Bound on Imbalance Now we can bound the total imbalance in a strip s at time τ'. Let $\kappa = \frac{e^2 \ln 2}{\log_2(4/3)}\sqrt{T}(\sqrt{T}+\log_2(2n^k))+\frac{e^{12}\ln 2}{4\log_2(4/3)}T\log_2(2n^k\sqrt{T})$. We divide the time interval $[0,\tau']$ into two periods: $[0,\tau'-\kappa]$ and $[\tau'-\kappa+1,\tau']$.

- In the first period, we bound each $|X(d,\tau)|$ and $|Y(d,\tau)|$ by $7.5n$ as no strip can have more than $7.5n$ agents on it by Lemma 3.6. By Claims 3.6.6 and 3.6.7, the total imbalance for this period is at most $15nTn^c/n^k$;
- For the second period, using Claims 3.6.3, 3.6.4, and 3.6.5, The total imbalance is at most $\lceil\kappa\rceil\cdot\left(192\sqrt{\frac{n\ln(4n^{3c+1}(T^2/32+T/8)N)}{\sqrt{T}}}+2\sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)}\right)$.

We choose $k=c+4$ and sum them up. We desire that both these contributions to the imbalance add up to no more than $n/25\sqrt{T}$. Using $n \geq T \geq 676$ (by the constraint (3.4)), we simplify this condition to conclude that it suffices to have:

$$n \geq (3654+2436e^{12}+546(e^{12}+1)c)^2(3c+4)T^3(\log_2 n)^2 \ln n.$$

The details of this calculation can be found in Appendix A.1.6.

Finally, the failure probability of $2/n^{2c+1}$ arises from Claims 3.6.4 and 3.6.5, which each have failure probability at most $1/n^{2c+1}$.

□

3.6.1.6 Initialization

Theorem 3.16. *Suppose that constraint* (3.4) *holds. If all agents follow the modified reasonable strategy, then* $H(\sqrt{T})$ *holds with probability at least* $1 - \frac{1}{n^{c+1}}$.

The proof is similar to the earlier analysis and can be found in Appendix A.1.7.

3.7 Numerical Simulations

We have demonstrated a strategy which is asymptotically close to optimal with regard to minimizing the average loss experienced by agents. Complementing this, in this section we simulate the evolution of the system for moderately large values of n and T. In order to gain a sense of the overall stability of the system, we track the total population over time.

We now discuss some observations based on our simulations.[4]

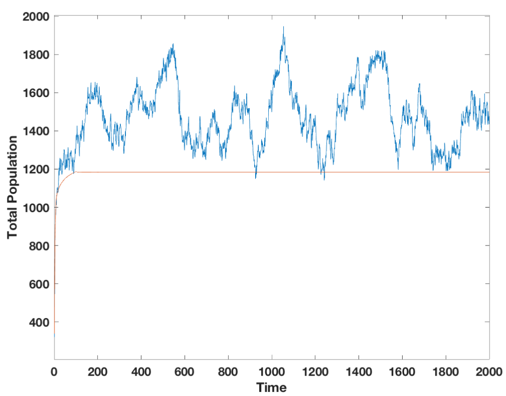

Figure 3.2: The evolution of total population over time in the discrete (blue) and continuum (orange) settings for $n = 500$, $T = 100$, using the reasonable strategy.

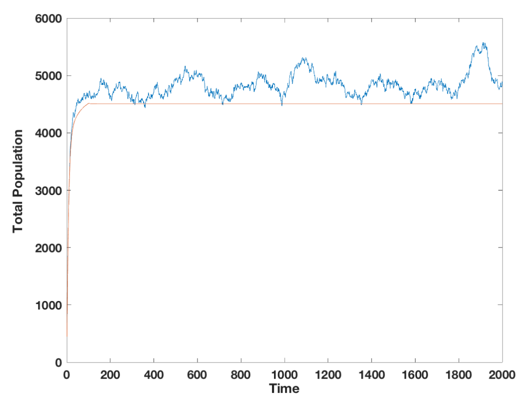

Figure 3.3: The evolution of total population over time in the discrete (blue) and continuum (orange) settings for $n = 500$, $T = 100$, using the modified reasonable strategy.

[4]For every pair of n and T that we considered in the discrete setting, we ran the simulation 10 times; letting each run for 2000 iterations. The error ranges mentioned below are obtained from the range of values we obtained over these 10 runs. The values for each run can be found in Appendix A.2.

For the continuum model we obtain reasonably rapid convergence—in about T time—whereas for the discrete model in a similar time the system reaches its long-term average value, but with somewhat chaotic oscillations about this value, as shown in Figures 3.2 and 3.3. In addition, the long-term average population for the discrete case is a bit larger than the continuum equilibrium value. (This is not surprising, for both variance and male/female imbalances will reduce the match rate.)

For moderate values of n and T, the average loss in the modified reasonable strategy is better than the asymptotic bound we obtain. For example, consider the $n = 500$, $T = 100$ case. We prove an upper bound on the total loss of $11T\sqrt{T}$ but the simulation achieves an average loss of just $2.21T\sqrt{T}(\pm 1.7\%)$. The total populations are significantly closer (an upper bound of close to $1.5nN$ in our theorem vs. close to $n\sqrt{T}$ in the simulation).

For the case where agents use the modified reasonable strategy, we also examine the average population size and average loss for various T (with n fixed at 500). Figure 3.4 shows a plot of Average Population/$n\sqrt{T}$ and Average Loss/$\sqrt{T}$ for five different values of n. The result is quite consistent with the $n\sqrt{T}$ scaling of the average total size and the $\sqrt{T}$ scaling of the average loss that we prove hold asymptotically, even though these are only moderately large values of n and T, and even though we are not in the n much greater than T regime of our analysis.

Finally, we examine the population on a "typical" diagonal[5] in the case that all agents are following the modified strip strategy (Figure 3.5). Notice that the range of oscillations for this value is large compared to the range of oscillations in the total population (Figure 3.3). Furthermore, these oscillations proceed at a much faster rate than the changes in the overall population. The results also indicate that while the system remains within reasonable bounds, there is substantial ongoing variation, particularly at a local level.

[5]Here we consider the diagonal region with width 1 that starts at value $3T/2$ at the top ($t = 0$) boundary.

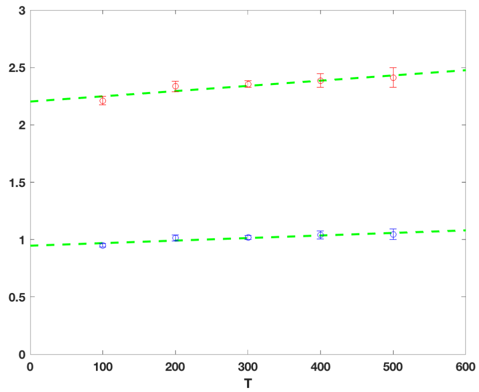

Figure 3.4: Average Population/$n\sqrt{T}$ (blue) and Average Loss/$\sqrt{T}$ (red) for five values of T, using the modified reasonable strategy.

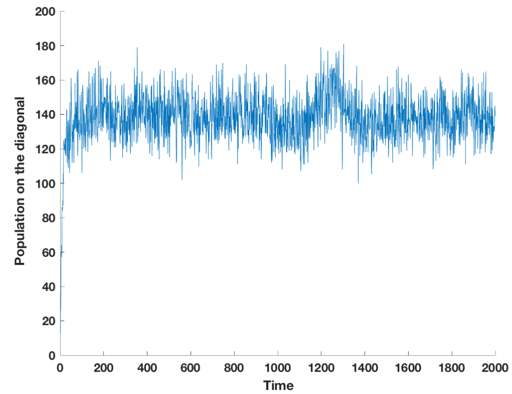

Figure 3.5: The evolution of the population on a "typical diagonal" over time in the discrete setting (for $n = 500$, $T = 100$), using the modified reasonable strategy.

3.8 Open Problems

Two natural extensions of our model come to mind.

- The values in our model are common to all agents, but in reality agents will have individual preferences. This could be captured with a model in which each agent a has a value for agent b given by $v_a + w_{a,b}$, where v_a is a common public value while $w_{a,b}$ is an idiosyncratic private value of a for b. The combining of public and private values has been studied in the literature on matchings in other settings [Ashlagi et al. 2019; Lee 2016].

- In our model, agents receive match proposals that are generated by choosing agents from the other side of the market uniformly at random. It would be interesting to consider a more sophisticated method of recommending matches, with recommended matches being localised in value and time around the agent.

Another intriguing direction concerns the stability of this system. We have shown that if the agents play the modified reasonable strategy then with high probability the strip sizes, the total size, and the imbalance between men and women in any strip, all remain within some range. But we conjecture that if any of these parameters were to have a large deviation which took it outside its typical range, then with high probability it would soon return to being within this range.

Finally, the Bayes-Nash Equilibrium strategy for the agents seems likely to be highly complex and dependent on the local populations throughout the system. Characterizing this equilibrium, as well as the loss bounds that can be achieved in such an equilibrium setting (or a simpler and perhaps more plausible ϵ-Bayes-Nash Equilibrium setting), remain open questions.

4 | Conclusion

As mentioned in Chapter 1, several mechanisms that have poor worst case guarantees seem to work quite well in practice. An underlying theme of this thesis has been to characterize situations where simple/commonly used matching algorithms appear to result in reasonable outcomes and to explain why they do so. We also sought to explain people's behaviour and/or provide useful advice in the form of simple strategies that agents can follow that have a high probability of yielding reasonably good matches.

In Chapter 2 we studied the stable matching problem in the public-private setting introduced by Lee [Lee 2016]. The two basic questions we sought to address are:

- Why does the deferred acceptance algorithm seem to work well even when agents submit short preference lists?
- Given that agents submit short preference lists, which prospective partners should they include?

We showed that, for a fairly general class of utility functions, in every stable matching, all but the bottom ϵ fraction of agents (as per their public rating) will be matched with an agent having a ranking (as per their public rating) close to their own. We considered a benchmark utility for each agent; we set the benchmark for an agent as the utility derived from matching with someone with the same ranking as per their public rating and the best possible private value. We showed that, in every stable matching, all but the bottom most agents obtain a match that yields utility close to their benchmark. We generalized these results to the many to one setting and also showed that our bounds are close to tight. Finally, we were able to demonstrate an ϵ-Bayes-Nash equilibrium where all but the bottommost agents make only $\Theta(\ln n)$ proposals and no agent makes more than $\Theta(\ln^2 n)$ proposals. Once again, in this approximate equilibrium, the agents propose to agents with similar public rating.

Our results point toward a rather intuitive answer to the question of which prospective partners an agent should apply to: it suffices to apply to their favourite prospective partners with a similar ranking as per public rating. This is because applying to a few prospective partners that are ranked about as highly as you are publicly perceived to be ranked suffices to obtain a utility which is close to the best you could achieve in any stable matching.

In Chapter 3 we studied a different setting; we considered a dynamic matching market where, at each time step, agents are paired up at random and can choose either to accept the pairing or to reject the pairing. A match occurs only if both agents choose to accept the pairing. We investigated how agents should behave when faced with the task of finding a match when they have to balance the trade-off between exploring for better matches and accepting a match now so as to enjoy a match for a longer period of time.

In our model, each agent has an individual value v_i and the same total lifetime T. The utility an agent obtains from a match is the value of their partner multiplied by the duration of the match, which is the minimum of their and their partner's remaining lifetimes. While this is quite a crude approximation of any real world setting, it still captures important qualitative features of agent behaviour: for example, as might be expected, it turns out that the agents should get less picky with time.

For an agent i we considered $v_i \cdot T$ as a benchmark for the utility. This corresponds to the utility the agent would derive if they were matched with an agent with the same value as themselves for the maximum possible duration. We provided a simple strategy, which if followed by all the agents, results in a low average loss for the agents, when the losses are computed with respect to the above defined benchmark. We also proved that our upper bound on the average loss is tight up to a constant factor. While our strategy yields low average loss if all agents follow it, characterizing a Nash equilibrium solution, ideally where the average loss is still low, remains open.

A | Appendices

A.1 Deferred Proofs from Section 3.6

Proof. (Of Lemma 3.6.) We complete the sketch proof by bounding the failure probability. Per time step, Theorems 3.9, 3.10, 3.12, 3.13 all have failure probability of at most $1/n^{2c+1}$, Theorem 3.15 has failure probability at most $2/n^{2c+1}$, and Theorem 3.16, which is applied once, has failure probability at most $1/n^{c+1}$. Theorem 3.14 does not introduce any additional possibility of failure. Multiplying by the n^c possible time steps, gives a total failure probability of at most $7/n^{c+1} < 1/n^c$. □

A.1.1 Upper Bound on Loss due to a Match

Proof. (of Lemma 3.7). Consider an agent (Agent 1) at value v and time t. Suppose they match with another agent (Agent 2) who is present in the same strip. The worst location for Agent 2 is to be on the low value strip boundary, and on this boundary to be at one of the endpoints.

Type 1 *strip.* If Agent 2 is at the top endpoint, Agent 1 obtains utility $w \cdot T - t$, where w is the value at the top endpoint. We can see that $w \geqslant v - \sqrt{T} - 2t$ (move from v horizontally to the lower boundary, a distance of at most $\sqrt{T}$ and then move up to the $t = 0$ location, which subtracts $2t$ from the value). Thus the utility Agent 1 receives is at least $(T-t)(v - \sqrt{T} - 2t)$. Therefore the loss is at most $t(v + 2T + \sqrt{T}) \leqslant 4tT + 2T\sqrt{T}$.

If Agent 2 is at the lower endpoint of a Type 1 strip, we argue as follows. The $v \cdot T - t$ product is equal at the two endpoints of a boundary, and therefore the loss is greatest at the top endpoint, for the utility garnered by Agent 1 would be $w \cdot (T - t)$ and not $w \cdot T$, whereas at the bottom endpoint the garnered utility is $2T \cdot w/2 = wT$.

Type 2 *Strip.* We define the following t values: a is the value for the left end of the top boundary of the strip, b the value for the left end of the bottom boundary. and c the value for the right end of the bottom boundary, Then $a \leqslant (2t + T - v)/2$, $b \leq 2a + \sqrt{T}$, and $c = T/2 + b \leq T/2 + 2a + \sqrt{T}$.

If Agent 2 is at the lower endpoint, then Agent 1 would receive utility $2T \cdot (T - c) \geqslant 2T(v - T/2 - 2t - \sqrt{T})$. Thus the loss is at most $T(T - v) + 4tT + 2T\sqrt{T} \leqslant 4tT + 2T\sqrt{T}$.

If Agent 2 is at the top endpoint and Agent 1 is older than Agent 2, then Agent 1 receives utility $T(T - t)$. As we are in a Type 2 strip, $(v - T) \leqslant 2t$ or $v \leqslant T + 2t$. So Agent 1 incurs a loss of at most $vT - T(T - t) \leqslant (T + 2t)T - T(T - t) \leqslant 3tT$.

While if Agent 1 is no older than Agent 2, then Agent 1 receives utility $T(T - b) \geqslant T(v - 2t - \sqrt{T})$. Thus the loss is at most $vT - T(v - 2t - \sqrt{T}) = 2Tt + T\sqrt{T}$. □

A.1.2 Lower Bound on the Total Population

Proof. (of Theorem 3.9) The agents enter with one of T values chosen uniformly at random and are equally likely to be men or women. Hence, for all n^c time steps, for each value v,

$$\Pr\left[\text{At most } \frac{n(1+\epsilon)}{2T} \text{ men enter with value } v\right] \geq 1 - n^c T e^{-\frac{\epsilon^2 n}{6T}}.$$

Call this event $\mathcal{E}$. Henceforth we condition on $\mathcal{E}$.

Let's consider those agents that enter at times in the range $[\tau - \sqrt{T} + 1, \tau]$ for some $\tau \leqslant n^c$. We want to lower bound the number of these agents who are present in the pool for the match at time τ.

In fact, henceforth, We will only consider men with values in the range $[T + \sqrt{T}, 2T)$. Among these men, consider those who have been in the pool for t time, where $0 \leqslant t < \sqrt{T}$. Let p_i^t be the probability that during their tth time step, the men in strip i are offered a match in their own strip. Even if all these men were still present in the matching pool,

$$\Pr\left[\text{\# of these men matched in strip } i \text{ at age } t \leqslant \frac{n(1+\delta)(1+\epsilon)\cdot p_i \cdot w_i}{2T}\right] \geq 1 - e^{-\frac{\delta^2 n p_i w_i}{6T}},$$

where w_i is the horizontal width of strip i occupied by these men when aged t. For every Type 1 strip, $w_i \leqslant \sqrt{T}$. For the one Type 2 strip, since all values are at least $T+\sqrt{T}$, for ages up to $\sqrt{T}$, $w_i \leqslant \sqrt{T}$. By applying $\overline{\mu} = \frac{n(1+\epsilon)\max\{p_i,\frac{1}{T}\}\sqrt{T}}{2T}$ in Lemma 3.5, it follows that:

$$\Pr\left[\text{\# of these men matched in strip } i \text{ at age } t \leqslant \frac{n(1+\delta)(1+\epsilon)\cdot \max\{p_i,\frac{1}{T}\}}{2\sqrt{T}}\right]$$
$$\geq 1 - e^{-\frac{\delta^2 n \max\{p_i,\frac{1}{T}\}\sqrt{T}}{6T}} \geq 1 - e^{-\frac{\delta^2 n}{6T^{1.5}}},$$

The sum of the match probabilities—the p_is— is at most 1. Notice that at any fixed time we only need to consider $\sqrt{T}$ strips, because at any time step, the men we are considering will occupy only $\sqrt{T}$ many strips. This implies $\sum \max\{p_i, \frac{1}{T}\} \leq 1 + \frac{1}{\sqrt{T}}$. Therefore,

$$\Pr\left[\begin{array}{l}\text{\# of these men being matched}\\ \text{over all the strips at age } t\end{array} \leqslant \frac{(1+\frac{1}{\sqrt{T}})n(1+\delta)(1+\epsilon)}{2\sqrt{T}}\right] \geq 1 - \sqrt{T}\cdot e^{-\frac{\delta^2 n}{6T^{1.5}}}.$$

Hence, we can bound the probability of the number of men who entered at time $\tau-\Delta+1$ and left by time τ, for any $\Delta \leqslant \sqrt{T}$, as follows:

$$\Pr\left[\begin{array}{l}\text{\# men being matched}\\ \text{in their first } \Delta \text{ steps}\end{array} \leqslant \frac{(1+\frac{1}{\sqrt{T}})n\Delta(1+\delta)(1+\epsilon)}{2\sqrt{T}}\right] \geq 1 - \Delta T\cdot e^{-\frac{\delta^2 n}{6T^{1.5}}}.$$

Consequently, we can bound the probability for the number of men that enter in the time interval $[\tau-\sqrt{T}+1, \tau-1]$ and are matched no later than time $\tau-1$ as follows, where we sum over all $\tau \leqslant n^c$:

$$\Pr\left[\begin{array}{l}\text{\# men who entered and were} \\ \text{matched in a } \sqrt{T}-1 \text{ time window}\end{array} \leqslant \frac{(1+\frac{1}{\sqrt{T}})n(1+\delta)(1+\epsilon)(\sqrt{T}-1)}{4}\right]$$
$$\geq 1 - \frac{1}{2} n^c T^{1.5} e^{-\frac{\delta^2 n}{6T^{1.5}}}.$$

We set $\delta = \left[\frac{6T^{1.5}}{n}\ln\left(10n^{2c+1}n^cT^{1.5}\right)\right]^{1/2}$ and $\epsilon = \left[\frac{6T}{n}\ln\left(20n^{2c+1}n^cT\right)\right]^{1/2}$. By constraint (3.4), $c \geq 1$, $400 \leq T \leq n$, and $n \geq 96T^2(3c+3)\ln n$, therefore $\delta \leq 1/4$ and $\epsilon \leq 1/64$. This yields the bound:

$$\Pr\left[\begin{array}{l}\text{\# men who entered in a } \sqrt{T}-1 \\ \text{window being matched}\end{array} \geqslant \frac{\frac{65}{64}\cdot\frac{5}{4}n\sqrt{T}}{4}\right] \geq 1 - \frac{1}{20n^{2c+1}}.$$

Since we have been conditioning on $\mathcal{E}$, this bound holds with probability at least $1 - \frac{1}{10n^{2c+1}}$. The same bound applies to the women.

Recalling that we excluded the men with values less than $T + \sqrt{T}$, this yields the following lower bound on the total population size, throughout this n^c time period:

$$n\sqrt{T} - \frac{n(1+\epsilon)}{2\sqrt{T}} - 0.635n\sqrt{T} \geq \frac{1}{3}n\sqrt{T},$$

with probability at least $1 - \frac{1}{5n^{2c+1}}$. □

A.1.3 Upper Bound on The Total Population

Proof. (Of Theorem 3.10.) Let $P(t)$ be the total population at the start of time step t. Let N be the total number of strips. By Constraint 3.4, $T \geq 676$, so $N \leq \sqrt{T} + \log_2\sqrt{T} + 1 \leq 5\sqrt{T}/4$.

If $P(t) \leqslant \frac{3}{2}nN$, then $P(t+1) \leqslant \frac{3}{2}nN + n$. So we will only consider the case that $P(t) > \frac{3}{2}nN$. In this case, the average strip population at the start of step t is more than $\frac{3}{2}n$.

Next, we upper bound the number of men in the population; the same bound applies to the number of women.

By $H(t)$, clause 5, the excess of men over women in each strip is at most $n/25\sqrt{T}$ except for the last Type 2 strip. So the excess over all these $N-1$ strips is at most $n(N-1)/25\sqrt{T}$. For the last Type 2 strip, the population is less than $60n/\sqrt{T}$ which is smaller than $40P(t)/676$ as $T \geq 676$. Consequently, there are at most $P(t)/2 + nN/50\sqrt{T} + 20P(t)/676 \leqslant 11P(t)/20$ men in the total population.

The expected number of matches in strip i, μ_i, is given by

$$E[\mu_i] = \frac{(\text{ number of women in strip } i) \times (\text{ number of men in strip } i)}{\text{number of men in the whole population}}.$$

Let s_i denote the population of the i-th strip. The denominator is at most $\frac{11}{20}P(t)$ and at least $\frac{1}{2}P(t)$. The numerator is minimized when the number of women and men in the strip are as far apart as possible. So, for the strips other than the last Type 2 strip, the numerator is at least $(s_i/2 + n/50\sqrt{T})(s_i/2 - n/50\sqrt{T}) = s_i^2/4 - n^2/2500T^2$. The numerator is maximized when the numbers of women and men are equal. Therefore,

$$\frac{s_i^2/4 - n^2/2500T^2}{\frac{11}{20}P(t)} \leq E[\mu_i] \leq \frac{s_i^2}{2P(t)}. \tag{A.1}$$

Consider an indicator random variable X_i for each man in this strip, which is 1 if that man gets matched. By Lemma 3.4 we can use a Chernoff bound to obtain:

$$\Pr\left[\text{number of matches in strip } i \leq \mu_i(1-\epsilon)\right] \leq e^{-\epsilon^2\mu_i/2}. \tag{A.2}$$

For $E[\mu_i] \geqslant \alpha n/\sqrt{T}$,

$$\Pr\left[\text{number of matches in strip } i \leq E[\mu_i](1-\epsilon)\right] \leq e^{-\epsilon^2 E[\mu_i]/2} \leq e^{-\epsilon^2\alpha n/(2\sqrt{T})}.$$

Let $\epsilon = \sqrt{\frac{2\sqrt{T}}{\alpha n}\ln(Nn^{2c+1})}$. Then $\Pr\left[\text{number of matches in strip } i \leq \mu_i(1-\epsilon)\right] \leq \frac{1}{Nn^{2c+1}}$

For $E[\mu_i] < \alpha n/\sqrt{T}$, let $\epsilon = \frac{\theta}{E[\mu_i]}$, then (A.2) becomes

$$\Pr\left[\text{number of matches in strip } i \leq E[\mu_i] - \theta\right] \leq e^{-\theta^2/(2E[\mu_i])} \leq e^{-\theta^2\sqrt{T}/(2\cdot\alpha n)}.$$

Let $\theta = \sqrt{\frac{2\alpha n}{\sqrt{T}}\ln(Nn^{2c+1})} = \alpha\epsilon\frac{n}{\sqrt{T}}$, then $\Pr\left[\text{number of matches in strip } i \leq \mu_i - \theta\right] \leq \frac{1}{Nn^{2c+1}}$

Let NL be the set of all strips except for the last Type 2 strip. Then, with probability at least $1 - \frac{1}{n^{2c+1}}$, the number of matches is larger than $(1-\epsilon)\sum_{i\in\mathsf{NL}} E[\mu_i] - N\theta$. In addition, by (A.1), $\sum_{i\in\mathsf{NL}} E[\mu_i]$ is lower bounded by:

$$\begin{aligned}\sum_{i\in\mathsf{NL}}\frac{s_i^2/4 - n^2/2500T^2}{\frac{11}{20}P(t)} &\geq \sum_{i\in\mathsf{NL}}\frac{\frac{1}{4}\left(\frac{9P(t)}{10N}\right)^2 - n^2/2500T^2}{\frac{11}{20}P(t)}\\ &\geq \left(\frac{81}{220}\frac{P(t)}{N} - \frac{n^2N}{1375T^2P(t)}\right)\\ &\geq \left(\frac{243}{440}n - \frac{2n}{4125T^2}\right)\\ &\geq \frac{11}{20}n.\end{aligned}$$

The first inequality follows as $\sum_{i\in\mathsf{NL}} s_i \geqslant \frac{9}{10}P(t)$, and the next to last inequality follows as $P(t) \geq \frac{3}{2}nN$. Let $\alpha = 0.4$. Since $\epsilon \leq \frac{1}{22}$, as $n \geq 2420\sqrt{T}(2c+2)\ln n$, and $N\theta \leq \frac{5}{4}\alpha\epsilon n$, as $N \leq \frac{5}{4}\sqrt{T}$ and $\theta = \alpha\epsilon n/\sqrt{T}$,

$$(1-\epsilon)\sum_{i\in\mathsf{NL}} E[\mu_i] - N\theta \geq (1-\epsilon)\frac{11}{20}n - \frac{5}{4}\alpha\epsilon n \geq \frac{n}{2}.$$

This means the total number of people matched in the market is greater than n, which is the number of people entering, which completes the proof.

$\square$

A.1.4 Upper Bound on the Size of a Type 1 Strip.

Proof. (Of Theorem 3.12) Consider any Type 1 strip s. For any two points (v, t_1) and (v, t_2) in s which have the same value v, we have $|t_2 - t_1| \leq \frac{\sqrt{T}}{2}$. Let s' be the strip immediately to the right of s.

We are going to lower bound the number of matches in time step t. Let's consider the agents who will be in strip s at time $t+1$. They will all enter s during a length $\sqrt{T}/2$ time interval ending at time $t+1$.

Let $P_{t-\sqrt{T}+1}$ be the agents in strip s' at time $t - \sqrt{T}$. By the inductive hypothesis, applied to s' at time $t - \sqrt{T}$, we know that $|P_{t-\sqrt{T}+1}| \leqslant dn$. We are going to track the subset of these agents who remain in the system after each step of matches, for the next $\sqrt{T}$ steps, along with the new agents who join the diagonals used by this subset of agents. Define $S_{t-\sqrt{T}+i}$ to be the rightmost $\sqrt{T}+1-i$ diagonals in s' plus the leftmost $i-1$ diagonals in s, for $1 \leqslant i \leqslant \sqrt{T}+1$. Let $P_{t-\sqrt{T}+i}$ be the population occupying $S_{t-\sqrt{T}+i}$ at the start of step $t - \sqrt{T} + i$. Then $P_{t-\sqrt{T}+i+1}$ is obtained from $P_{t-\sqrt{T}+i}$ by removed matched agents and then adding the new agents for the diagonals in $S_{t-\sqrt{T}+i+1}$. Our analysis will show that, with high probability, for each of these $\sqrt{T}$ steps, the number of matches is at least the number of new agents. This implies the upper bound on the strip population continues to hold.

By means of a Chernoff bound, we observe that the number of new agents per step can be bounded with high probability as follows:

$$\Pr\left[\#\text{ new agents } \leq \frac{n}{\sqrt{T}}(1+\delta)\right] \geqslant 1 - e^{-\frac{n\delta^2}{3\sqrt{T}}}. \tag{A.3}$$

Let $\delta = \sqrt{\frac{3\sqrt{T}}{n}\ln(Nn^{2c+1})}$. As $n \geqslant 60T(2c+2)\ln n$, $\delta \leq \frac{1}{20}$, which yields

$$\Pr\left[\#\text{ new agents} \leq \frac{41n}{40\sqrt{T}} = 1.025\frac{n}{\sqrt{T}}\right] \geqslant 1 - \frac{1}{Nn^{2c+1}}. \tag{A.4}$$

By (3.6), the maximum of the number of men and number of women in the market is at most $n\sqrt{T}$.

Let $P_{t'',s}$ and $P_{t'',s'}$ be the portions of population $P_{t''}$ at time t'' in strips s and s', resp., for $t + 1 - \sqrt{T} \leqslant t'' \leqslant t$.

By Lemma 3.11, the matches remove at least the following number of people from $P_{t'',s}$:

$$\frac{(|P_{t'',s}|^2)/2 - (n/25\sqrt{T})^2/2}{n}\sqrt{T} = \frac{|P_{t'',s}|^2 - n^2/625T}{2n\sqrt{T}}$$

A similar bound applies to the matches involving $P_{t'',s'}$. To minimize the terms $|P_{t'',s'}|^2/2n\sqrt{T}$ for $P_{t'',s}$ and $P_{t'',s'}$, we should make them equal. Thus the expected number of matches of population $P_{t''}$ is at least

$$\frac{|P_{t''}|^2}{4n\sqrt{T}} - \frac{n}{625T\sqrt{T}}. \tag{A.5}$$

Next, we want to obtain a high probability bound.

There are four sets of people, resp. the men and women in each of $P_{t'',s}$ and $P_{t'',s'}$. Let μ be the number of matches of one set. If $E[\mu] \geq \frac{\alpha n}{\sqrt{T}}$, then

$$\Pr\left[\mu \geqslant E[\mu](1-\epsilon)\right] \geqslant 1 - e^{-(E[\mu]\epsilon^2)/2} \geqslant 1 - e^{-\epsilon^2 \alpha n/(2\sqrt{T})};$$

letting $\epsilon = \sqrt{\frac{2\sqrt{T}}{\alpha n}\ln(4Tn^{2c+1})}$ yields $\Pr\left[\mu \geqslant E[\mu](1-\epsilon)\right] \geqslant 1 - \frac{1}{4Tn^{2c+1}}$.

Otherwise, $E[\mu] \leq \frac{\alpha n}{\sqrt{T}}$, and

$$\Pr\left[\mu \geqslant E[\mu] - \theta\right] \geqslant 1 - e^{-(\theta^2)/(2E[\mu])} \geqslant 1 - e^{-(\sqrt{T}\theta^2)/(2\alpha n)};$$

setting $\theta = \sqrt{\frac{2\alpha n}{\sqrt{T}}\ln\left(4\sqrt{T}n^{2c+1}\right)} = \alpha\epsilon\frac{n}{\sqrt{T}}$ yields $\Pr\left[\mu \geqslant E[\mu] - \theta\right] \geqslant 1 - \frac{1}{4Tn^{2c+1}}$.

Recall (A.4), the high probability bound that the number of new agents is at most $1.025\frac{n}{\sqrt{T}}$.

By (A.5), the number of people matched is at least $(1-\epsilon)\left[\frac{|P_{t''}|^2}{4n\sqrt{T}} - \frac{n}{625T\sqrt{T}}\right] - 4\theta$. Recall that $|P_{t''}| \le dn$; we let $x = dn - |P_{t''}|$. Then, the number of people left is at most:

$$dn - x - \left[(1-\epsilon)\left[\frac{(dn-x)^2}{3n\sqrt{T}} - \frac{n}{625T\sqrt{T}}\right] - 4\theta\right] \le dn - \left[(1-\epsilon)\left[\frac{(dn)^2}{4n\sqrt{T}} - \frac{n}{625T\sqrt{T}}\right] - 4\theta\right],$$

if $1 \ge (1-\epsilon)d/2\sqrt{T}$. This number is upper bounded by $dn - (1-\epsilon)(\frac{d^2}{4} - \frac{1}{422{,}500})\frac{n}{\sqrt{T}} + 4\alpha\epsilon\frac{n}{\sqrt{T}}$ as $T \ge 676$ and $|P_{t''}| \ge dn$. Let $d = 2.6$ and $\alpha = \frac{3}{16}$. Also, $\epsilon \le \frac{1}{10}$ as $n \geqslant 27(2c+2)T\ln n$ and $676 \leqslant T \le n$. A final calculation shows that $(1-\epsilon)(\frac{d^2}{4} - \frac{1}{422{,}500})\frac{n}{\sqrt{T}} - 4\alpha\epsilon\frac{n}{\sqrt{T}}$ is at least $1.025\frac{n}{\sqrt{T}}$, demonstrating the result.

□

A.1.5 Upper Bound on the Size of a Type 2 Strip

Proof. (of Theorem 3.13)

Consider any Type 2 strip s. If s is the topmost Type 2 strip, clearly we can upper bound its size by twice the bound on the size of a Type 1 strip given in Theorem 3.12. In addition, if s is the Type 2 strip next to the topmost Type 2 strip, then the size of strip s is less than that of the topmost Type 2 strip at time $t + 1 - \sqrt{T}$, which completes the proof for this strip too.

We now assume that s has at least two Type 2 strips above it. Let s' be the strip immediately above s, and let h denote the height of s. Then the height of s' is $h/2$ and $h \ge \sqrt{T}$.

Let's consider the agents who will be in strip s at time $t+1$. They will all enter s during a length h time interval ending at time $t+1$. They can be partitioned into two sets as follows:

- Y_{t+1} = {agents that will have spent less than $h/2$ time in strip s by time $t+1$}. .
- O_{t+1} = {agents that will have spent at least $h/2$ time steps in strip s by time $t+1$}.

The agents in Y_{t+1} were all present at time $t' = t + 1 - h/2$ as part of the population of strip

s' at that time. By the inductive hypothesis, applied to s' at time t', we know that there were at most $2gn\sqrt{T}/h$ agents in s' at that time. Let P_y denote this population. The agents in O_{t+1} were all present at time $t' = t+1-h$ as part of the population of strip s' at that time. By the inductive hypothesis, applied to s' at time t', we know that there were at most $2gn\sqrt{T}/h$ agents in s' at that time. Let P_o denote this population.

Let $P^y_{t'',s}$ and $P^y_{t'',s'}$ be the remainder of population P_y at time t'' in strips s and s', resp., for $t+1-\sqrt{T} \leqslant t'' \leqslant t$. Also, let $P^y_{t''} = P_{t'',s} \cup P_{t'',s'}$. Similarly, define $P^o_{t'',s}$, $P^o_{t'',s'}$ and $P^o_{t''}$.

To this end, we need to compute lower bounds on the match rates.

By (3.6), the maximum of the number of men and number of women in the market is at most $n\sqrt{T}$.

We divide the period $[t+1-h, t+1)$ into two phases; Phase 1, $[t+1-h, t+1-h/2)$, and Phase 2, $[t+1-h/2, t+1)$. We will show that the size of $P^y_{t''}$ at the end of Phase 1 is at most $g_1 n\sqrt{T}/h$. We will specify g_1 later. Then, at the start of Phase 2 the size of $P^o \cup P^y$ is at most $2gn\sqrt{T}/h + g_1 n\sqrt{T}/h$. We claim that after Phase 2, the size of $P^o \cup P^y$ is at most $gn\sqrt{T}/h$.

We analyze Phase 1 first. Consider the set $P^o_{t'',s}$ and time $t'' \in [t+1-h, t+1-h/2)$.

By Lemma 3.11, at time t'' these matches remove, in expectation, at least the following number of people from $P^o_{t'',s}$:

$$\frac{|P^o_{t'',s}|^2 - n^2/625T}{2n\sqrt{T}} \tag{A.6}$$

Similar bounds will hold for the sets $P^o_{t'',s'}$. Notice that the total expected number of matches from $P^o_{t''}$ is minimized if $|P^o_{t'',s}| = |P^o_{t'',s'}|$. Thus we obtain that the size of $P^o_{t''}$ reduces, in expectation, by at least

$$\frac{|P^o_{t''}|^2}{4n\sqrt{T}} - \frac{n}{625T\sqrt{T}}. \tag{A.7}$$

As in the analysis for the Type 1 strip, we then give a high probability bound. We have four sets of people, the men and the women in the sets $P^o_{t'',s'}$ and $P^o_{t'',s}$, respectively. Suppose μ be the number of matches in one of these set at time t''.

If $E[\mu] \geq \alpha n\sqrt{T}/h^2$, by Lemma 3.4,

$$\Pr\left[\mu \geqslant E[\mu](1-\epsilon)\right] \geqslant 1 - e^{-(E[\mu]\epsilon^2)/2} \geqslant 1 - e^{-\epsilon^2 \alpha n\sqrt{T}/2h^2},$$

Setting $\epsilon = \left[\frac{2h^2}{\alpha n\sqrt{T}} \ln\left(T(\log_2 \sqrt{T}+1)n^{2c+1}\right)\right]^{\frac{1}{2}}$ yields $\Pr\left[\mu \geqslant E[\mu](1-\epsilon)\right] \geqslant 1 - \frac{1}{T(\log_2 \sqrt{T}+1)n^{2c+1}}$. Otherwise, $E[\mu] \leq \alpha n\sqrt{T}/h^2$, so by Lemma 3.4,

$$\Pr\left[\mu \geqslant E[\mu] - \theta\right] \geqslant 1 - e^{-(\theta^2)/(2E[\mu])} \geqslant 1 - e^{-(\theta^2 h^2)/(2\alpha n\sqrt{T})},$$

Setting $\theta = \left[\frac{2\alpha n\sqrt{T}}{h^2} \ln\left(T(\log_2 \sqrt{T}+1)n^{2c+1}\right)\right]^{\frac{1}{2}} = \epsilon\alpha\frac{n\sqrt{T}}{h^2}$ yields $\Pr\left[\mu \geqslant E[\mu] - \theta)\right] \geqslant 1 - \frac{1}{T(\log_2 \sqrt{T}+1)n^{2c+1}}$.

For each of the four sets we can use one of the two bounds above.

We can set $\alpha = 0.1, \epsilon \leq 0.1$ by imposing the constraints $c \geq 1, 400 \leq T \leq n, n \geq 125(2c+2.5)T\sqrt{T}\ln n$ which are provided by the constraints in (3.4), $h \leq T/4$, and $\log_2 \sqrt{T}+1 \leq \sqrt{T}/4$.

Let $g(\cdot)$ be a real valued function. Suppose the size of the set $P^o_{t''}$ at round t'' is smaller than $g(t'') \cdot n\sqrt{T}/h$ and let $X = g(t'') \cdot n\sqrt{T}/h - |P^o_{t''}|$. If $(1-\epsilon)g(t'')/(2h) \leq 1$, [1] then the size of the set $P^o_{t''+1}$ at round $t''+1$ is at most

$$\begin{aligned}
&g(t'')n\sqrt{T}/h - X - (1-\epsilon)\left[\frac{(g(t'')n\sqrt{T}/h - X)^2}{4n\sqrt{T}} - \frac{n}{625T\sqrt{T}}\right] + 4\epsilon\alpha\frac{n\sqrt{T}}{h^2} \\
&\leq g(t'')n\sqrt{T}/h - (1-\epsilon)\left[\frac{(g(t'')n\sqrt{T}/h)^2}{4n\sqrt{T}} - \frac{n}{625T\sqrt{T}}\right] + 4\epsilon\alpha\frac{n\sqrt{T}}{h^2} \\
&\leq g(t'')n\sqrt{T}/h - \frac{9g(t'')^2 n\sqrt{T}}{40h^2} + \frac{9n}{6250T\sqrt{T}} + \frac{n\sqrt{T}}{25h^2} \\
&\leq n\sqrt{T}/h\left[g(t'') - \frac{1}{h}\left(\frac{9g(t'')^2}{40} - 0.041\right)\right].
\end{aligned}$$

[1]Note that this is satisfied when $g(t'') \leq 21.5$ and $h \geq \sqrt{T} = 20$.

The last inequality uses the constraint that $h \le T/4$.

Let $g(t+1-h) = 2g$ and $g(t''+1) = \left[g(t'') - \frac{1}{h}\left(\frac{9g(t'')^2}{40} - 0.041\right)\right]$ for $t'' \in [t+1-h, t+1-h/2)$, then we have shown that the size of the set $P^o_{t''}$ at round $t+1-h/2$ is at most $g(t+1-h/2)\cdot n\sqrt{T}/h$. One way to solve $g(\cdot)$ by using a differential equation. Consider the differential equation $\mathrm{d}\bar{g}/\mathrm{d}t = -\frac{1}{h}\left(\frac{9\bar{g}^2}{40} - 0.041\right)$ and $\bar{g}(t+1-h) = 2g$. Note that $\bar{g}(t'') \ge g(t'')$ for all $t'' \in [t+1-h, t+1-h/2]$. [2]

Therefore, in order to prove $g(t+1-h/2) \le g_1$, we only need $\bar{g}(t+1-h/2) \le g_1$. We look at the total time for $\bar{g}$ to reduce from the value $g(t+1-h) = 2g$ to g_1: $\mathrm{d}t = -h\mathrm{d}\bar{g}/(\frac{9\bar{g}^2}{40} - 0.041)$. Therefore, the total time is $\int_{\bar{g}=g_1}^{2g} h\mathrm{d}\bar{g}/(\frac{9\bar{g}^2}{40} - 0.041) <= \int_{\bar{g}=g_1}^{2g} h\mathrm{d}\bar{g}/(\frac{9\bar{g}^2}{40} - 0.041\frac{\bar{g}^2}{g_1^2}) = (h/(\frac{9}{40} - \frac{0.041}{g_1^2}))(1/g_1 - 1/(2g))$. To have this be at most $h/2$ (the total duration of Phase 1), we only need $2(1/g_1 - 1/(2g)) \le 9/40 - 0.041/(g_1)^2$, which is satisfied by letting $g = 7.5$ and $g_1 = 6.5$.

We consider Phase 2 next. Consider the set $P^o_{t'',s} \cup P^y_{t'',s}$ and time $t'' \in [t+1-h/2, t+1)$. The analysis is exactly the same as that for Phase 1. By Lemma 3.11, these matches remove, in expectation, at least the following number of people from $P^o_{t'',s} \cup P^y_{t'',s}$:

$$\frac{|P^o_{t'',s} \cup P^y_{t'',s}|^2 - n^2/625T}{2n\sqrt{T}} \tag{A.8}$$

Similar bounds will hold for the set $P^o_{t'',s'} \cup P^y_{t'',s'}$. We also reduce the size of $P^o_{t''} \cup P^y_{t''}$, in expectation, by at least

$$\frac{|P^o_{t''} \cup P^y_{t''}|^2}{3n\sqrt{T}} - \frac{n}{625T\sqrt{T}}. \tag{A.9}$$

Then, as in Phase 1, suppose the size of the set $P^o_{t''} \cup P^y_{t''}$ at time t'' is smaller than $g(t'')n\sqrt{T}/h$ and let $X = g(t'')n\sqrt{T}/h - |P^o_{t''} \cup P^y_{t''}|$. If $(1-\epsilon)g(t'')/(2h) \le 1$, then the size of the set $P^o_{t''} \cup P^y_{t''}$

[2] Suppose it is not true. Since $\bar{g}(t+1-h) = g(t+1-h) = 2g$, there exists a t', such that $\bar{g}(t') \ge g(t')$ and $\bar{g}(t'+1) < g(t'+1)$. Then, there exist a $t'' \in [t', t'+1)$ such that $\bar{g}(t'') = g(t')$. After time t'', $\mathrm{d}\bar{g}(t'')/\mathrm{d}t \ge -\frac{1}{h}\left(\frac{9g(\bar{t}'')^2}{40} - 0.041\right) = g(t'+1) - g(t')$. Therefore, $\bar{g}(t'+1) = \bar{g}(t'') + \int_{s=t''}^{t'+1} \mathrm{d}\bar{g}(s) \ge g(t') + \int_{s=t''}^{t'+1} [g(t'+1) - g(t')]\mathrm{d}s \ge g(t'+1)$, which contradicts the assumption.

at round $t''+1$ is at most

$$n\sqrt{T}/h\left[g(t'') - \frac{1}{h}\left(\frac{9g(t'')^2}{40} - 0.041\right)\right].$$

Let $g(t+1-h/2) = 2g+g_1$ and $g(t''+1) = \left[g(t'') - \frac{1}{h}\left(\frac{9g(t'')^2}{40} - 0.041\right)\right]$ for $t'' \in [t+1-h/2, t+1)$, then we have shown that the size of the set $P^o_{t''} \cup P^y_{t''}$ at round $t+1$ is at most $g(t+1)\cdot n\sqrt{T}/h$. We consider the same differential equation here, $\mathrm{d}\bar{g}/\mathrm{d}t = -\frac{1}{h}\left(\frac{9\bar{g}^2}{40} - 0.041\right)$, with $\bar{g}(t+1-h/2) = 2g+g_1$. Note that $\bar{g}(t'') \geq g(t'')$ for all $t'' \in [t+1-h/2, t+1]$.

Therefore, in order to prove $g(t+1) \leq g$, we only need $\bar{g}(t+1) \leq g$. We look at the total time for $\bar{g}$ to reduce from the value $2g+g_1$ to g: $\mathrm{d}t = -h\mathrm{d}\bar{g}/(\frac{9\bar{g}^2}{40} - 0.041)$. Therefore, the total time is $\int_{\bar{g}=g}^{2g+g_1} h\cdot \mathrm{d}\bar{g}/(\frac{9\bar{g}^2}{40} - 0.041) <= \int_{\bar{g}=g}^{2g+g_1} h\cdot \mathrm{d}\bar{g}/(\frac{9\bar{g}^2}{40} - 0.041\frac{\bar{g}^2}{g^2}) = (h/(\frac{9}{40} - \frac{0.041}{g^2}))(1/g - 1/(2g+g_1))$. To have this be at most $h/2$ (the total duration of Phase 2), we only need $2(1/g - 1/(2g+g_1)) \leq 9/40 - 0.041/(g)^2$, which is also satisfied by letting $g = 7.5$ and $g_1 = 6.5$.

Finally, as there are $\log_2\sqrt{T}+1$ Type 2 strips, and, for each Type 2 strip, we consider $h \leq T/4$ steps, the success probability is at least $1 - \frac{T/4\cdot 4\cdot(\log_2\sqrt{T}+1)}{T(\log_2\sqrt{T}+1)n^{2c+1}} \leq 1 - \frac{1}{n^{2c+1}}$.

□

Proof. (Of Theorem 3.14.) Let's call the bottommost Type 2 strip s and the Type 2 strip immediately above it strip s'. Any agent in the population in s at time $t+1$ must belong to one of the following categories:

- The agent was present in s' at time $t+1-T/4$.
- The agent was present in s' at time $t+1-T/2$.

However, by our inductive hypothesis $H(t)$, we know that at all time steps before $t+1$, the size of s was always less than $\frac{7.5n\sqrt{T}}{\text{height of strip } s} \leq 4\cdot 7.5n/\sqrt{T}$.

Thus, the size of s at time $t+1$ is bounded by $8\cdot 7.5n/\sqrt{T} \leq 60n/\sqrt{T}$, which concludes the proof. □

A.1.6 Bound on the Imbalance

Proof. (Of Claim 3.6.1). The expected number of matches at time τ between men in diagonal d_i and women in diagonal d_j is

$$\frac{(2A_i + I_i + X_i)(2A_j - I_j - X_j)}{4R}.$$

Similarly, the expected number of women in d_i that match with men in d_j is

$$\frac{(2A_i - I_i - X_i)(2A_j + I_j + X_j)}{4R}.$$

Thus $I'(d_i, \tau)$ is given by:

$$\begin{aligned}
&I_i + X_i - \sum_{d_j \in s} \left[\frac{(2A_i + I_i + X_i)(2A_j - I_j - X_j)}{4R} - \frac{(2A_i - I_i - X_i)(2A_j + I_j + X_j)}{4R} \right] \\
&\quad = I_i + X_i - \sum_{d_j \in s} \left[X_i \frac{(2A_j - I_j - X_j)}{4R} - X_j \frac{(2A_i - I_i - X_i)}{4R} + I_i \frac{(2A_j - I_j - X_j)}{4R} - I_j \frac{(2A_i - I_i - X_i)}{4R} \right].
\end{aligned}$$

□

Proof. (Of Claim 3.6.5.) We first give a high probability bound on $\sum Y(d_i, \tau)$. Let $m(d_i, \tau)$ be the number of men entering the market on diagonal d_i at time τ. Note that d_T is the last Type 1 diagonal. Let d_r be a diagonal in Type 1 strip s; then,

$$\Pr\left[\left| \sum_{r \le i \le T} Y(d_i, \tau) \right| > \Delta \right] = \Pr\left[\left| \sum_{r \le i \le T} m(d_i, \tau) - (T - r + 1) \cdot \frac{n}{2T} \right| > \Delta/2 \right].$$

Note that

$$\Pr\left[\Big|\sum_{r\leq i\leq T} m(d_i,\tau) - (T-r+1)\cdot\frac{n}{2T}\Big| > \Delta/2\right]$$
$$\leq 2\exp\left[-\Delta^2/\Big(3\cdot(T-r+1)\cdot\frac{n}{2T}\Big)\right] \leq 2\exp\left[-\Delta^2/\Big(\frac{3n}{2}\Big)\right].$$

The last inequality follows as $T-r+1\leq T$. Letting $\Delta = \sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)}$ yields

$$\Pr\left[\Big|\sum_{r\leq i\leq T} Y(d_i,\tau)\Big| > \sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)}\right] \leq \frac{1}{Tn^{3c+1}}.$$

Therefore, with probability at least $\frac{1}{n^{3c+1}}$, for all r such that d_r is a diagonal in a Type 1 strip, $\left|\sum_{r\leq i\leq T} Y(d_i,\tau)\right| \leq \sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)}$.

With this result in hand, we prove the claim as follows. Let $d_{r(s)}$ be the rightmost (lowest index) diagonal in s and $d_{l(s)}$ be the leftmost (highest index) diagonal in s.

Let $w_i = \sum_{j\geqslant r(s)} Y(d_i,\tau,d_j,\tau')/Y(d_i,\tau)$. Let's consider $\sum_{d_i\in s';j\geqslant r(s)} Y(d_i,\tau,d_j,\tau') = \sum_{d_i\in s'} w_i\cdot Y(d_i,\tau)$. By Claim 3.6.2, $w_i \leqslant w_k$, for $i<k$. Thus,

$$\Big|\sum_{d_i;j\geqslant r(s)} Y(d_i,\tau,d_j,\tau')\Big| = \Big|\sum_{i=1}^{T} w_i\cdot Y(d_i,\tau)\Big| \leq w_T\max_{r\leq T}\Big|\sum_{r\leqslant i\leqslant T} Y(d_i,\tau)\Big| \leq \sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)}.$$

Finally,

$$\Big|\sum_{d_i;d_j\in s} Y(d_i,\tau,d_j,\tau')\Big| = \Big|\sum_{d_i;j\geqslant r(s)} Y(d_i,\tau,d_j,\tau') - \sum_{d_i;j\geqslant l(s)+1} Y(d_i,\tau,d_j,\tau')\Big| \leqslant 2\sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)},$$

which completes the proof. □

Proof. (Of Claim 3.6.7.) We begin by bounding $\sum_{j:d_j \in s}(2A_j - I_j - X_j)/4R$ for any Type 2 strip s.

$$\sum_{j:d_j \in s} \frac{(2A_j - I_j - X_j)}{4R} \le \frac{1}{2} \frac{\frac{3.75n\sqrt{T}}{H} + n/50\sqrt{T}}{\frac{n\sqrt{T}}{6}} < \frac{23}{2H} \quad (\text{as } \sqrt{T} \geqslant 26 \text{ by constraint 3.4}). \tag{A.10}$$

Let s have height H.

Consider $X(d, \tau, d', \tau')$. If d' is in a Type 2 strip then by (A.10) at most $\frac{23}{2H}$ of it disperses to some location in the same strip and at least $1 - \frac{23}{2H}$ of it moves down distance one. This implies that, within H time, a Type 2 strip loses at least $e^{-23/2}$ of the $X(d, \tau, d', \tau')$ that had been present within it at time τ'. Let $K_2 = e^{12} \ln 2$. Therefore, by time $\tau' + K_2 H$ ($\le \tau' K_2 T/4$) at least half of the $X(d, \tau, d', \tau')$ in a Type 2 strip has moved out of the strip.

Similarly, we can now carry out the same kind of argument for the Type 2 strips. After $2e^2(\ln 2)\sqrt{T}(\sqrt{T} + \log_2(2n^k))$ time there is at most $\frac{1}{2n^k}$ fraction of $X(d, \tau)$ in the Type 1 strips. We focus on the remaining $1 - \frac{1}{2n^k}$ portion of $X(d, \tau)$ which has already entering the type 2 strips. Number the Type 2 strips from top to bottom[3]. Now consider γ as a distribution of the rest of the $X(d, \tau)$ where γ_i is the fraction of $X(d, \tau)$ in strip number i. Recall that there are $\log_2(\sqrt{T})$ Type 2 strips other than the bottom Type 2 strip. We consider the worst case where all the remaining $(1 - \frac{1}{2n^k}) \cdot X(d, \tau)$ starts out at the topmost Type 2 strip. Define a potential function $\phi(\gamma) = \sum_{i=1}^{\log_2 \sqrt{T}+1} \gamma_i \cdot 2^{(\log_2 \sqrt{T}) - i + 1}$. For the remaining $(1 - \frac{1}{2n^k}) \cdot X(d, \tau)$, The initial potential is at most $\sqrt{T}$. Every $K_2 T/4$ time steps, the potential decreases by at least $1/4$. Therefore, after $\frac{1}{\log_2(4/3)} K_2 T/4 \log_2(2n^k\sqrt{T})$ time steps, the potential would have reduced to at most $\frac{1}{2n^k}$.

There is also $\frac{1}{2n^k}$ fraction which might still be in the Type 1 strips. Thus the fraction of $X(d, \tau)$ in any strip other than the bottommost Type 2 strip after $\frac{1}{\log_2(4/3)} K_1 \sqrt{T}(\sqrt{T} + \log_2(2n^k)) +$ $\frac{1}{\log_2(4/3)} K_2 T/4 \log_2(2n^k\sqrt{T})$ time is at most $\frac{1}{n^k}$. ◻

Proof. (Final Calculation in Theorem 3.15)

[3] Our argument doesn't involve the last Type 2 strip, so we will end at the second to last strip.

As $\kappa = \frac{e^2 \ln 2}{\log_2(4/3)}\sqrt{T}(\sqrt{T}+\log_2(2n^k))+\frac{e^{12}\ln 2}{4\log_2(4/3)}T\log_2(2n^k\sqrt{T}) \le 12.35(T+\sqrt{T}+(c+4)\sqrt{T}\log_2 n)+\frac{e^{12}}{2}(T+(c+4)T\log_2 n+0.5T\log_2 T)$, the total imbalance is at most

$$\frac{15T}{n^3}+\left[12.35(T+\sqrt{T}+(c+4)\sqrt{T}\log_2 n)+\frac{e^{12}}{2}(T+(c+4)T\log_2 n+0.5T\log_2 T)\right]\cdot \left(192\sqrt{\frac{n\ln(4n^{3c+1}(T^2/32+T/8)N)}{\sqrt{T}}}+2\sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)}\right).$$

In order to make it smaller than $\frac{n}{25\sqrt{T}}$, we only need that

$$\frac{375T\sqrt{T}}{n^3}+\left[309(T+\sqrt{T}+(c+4)\sqrt{T}\log_2 n)+\frac{25e^{12}}{2}(T+(c+4)T\log_2 n+0.5T\log_2 T)\right]\cdot \left(192\sqrt{\sqrt{T}\ln[4n^{3c+1}(T^2/32+T/8)N]}+2\sqrt{\frac{3T}{2}\ln\left(2Tn^{3c+1}\right)}\right) \le \sqrt{n}.$$

As $375T\sqrt{T}/n^3 \le 0.0012\sqrt{n}$ from the constraint $n \ge T \ge 676$, we need

$$\left[309(T+\sqrt{T}+(c+4)\sqrt{T}\log_2 n)+\frac{25e^{12}}{2}(T+(c+4)T\log_2 n+0.5T\log_2 T)\right]\cdot \left(192\sqrt{\sqrt{T}\ln[4n^{3c+1}(T^2/32+T/8)N]}+2\sqrt{\frac{3T}{2}\ln\left(2Tn^{3c+1}\right)}\right) \le 0.9988\sqrt{n}. \tag{A.11}$$

In addition, as $n \ge T \ge 676$,

$$\begin{aligned}&\left[309(T+\sqrt{T}+(c+4)\sqrt{T}\log_2 n)+\frac{25e^{12}}{2}(T+(c+4)T\log_2 n+0.5T\log_2 T)\right]\\ &\qquad \le (86.61+12.876c+57.62e^{12}+12.5e^{12}c)T\log_2 n,\end{aligned} \tag{A.12}$$

and, as $n \ge T \ge 676$ and $n \ge N$,

$$\left(192\sqrt{\sqrt{T}\ln[4n^{3c+1}(T^2/32+T/8)N]}+2\sqrt{\frac{3T}{2}\ln\left(2Tn^{3c+1}\right)}\right) \le 42\sqrt{T(3c+4)\ln n}. \tag{A.13}$$

By inequalities (A.11), (A.12), and (A.13),

$$(86.61 + 12.876c + 57.62e^{12} + 12.5e^{12}c)T \log_2 n \cdot 42\sqrt{T(3c+4)\ln n} \leq 0.9988\sqrt{n}.$$

Therefore, $n \geq (3654 + 2436e^{12} + 546(e^{12}+1)c)^2(3c+4)T^3(\log_2 n)^2 \ln n$ suffices.

□

A.1.7 Initialization

Proof. (Of Theorem 3.16.) At any point in the first $\sqrt{T}$ time steps:

- The total population in the entire matching pool is clearly less than $n\sqrt{T} < nN$, as only these many agents could have even entered the matching pool.
- In any single Type 1 strip,

  $$\Pr\left[\text{Number of agents that entered the strip from the top} \leq \frac{n\sqrt{T}(1+\epsilon)}{\sqrt{T}}\right] \geq 1 - e^{-\frac{n\epsilon^2}{3}}.$$

  However the agents that enter a Type 1 strip during the first $\sqrt{T}$ time must either have entered from the top or they could have entered at the top boundary of the previous strip. Thus, by a union bound,

  $$\Pr\left[\text{Number of agents that entered any Type 1 strip} \leq 2n(1+\epsilon)\right] \geq 1 - \sqrt{T}e^{-\frac{n\epsilon^2}{3}}.$$

  So by setting $\epsilon = \sqrt{\frac{3}{n}\ln\left(n^{c+1}\sqrt{T}\right)}$, and imposing the constraints $c \geqslant 1$, $T \leqslant n$, and $n \geqslant 35(c+2)\ln n$ that guarantee that $\epsilon < 0.3$ (from (3.4)), we obtain that with probability $1 - \frac{1}{n^{c+1}}$ every Type 1 strip has a population $< 2.6n$.

- The agents in the first Type 2 strip after $\sqrt{T}$ time steps (the only Type 2 strip with any population after $\sqrt{T}$ time) could only be those agents that entered the leftmost two Type 1 strips from the top. However the previous bound already guarantees that this number is also $< 2.6n$.

- Also, the population in the bottom most Type 2 strip will be 0.

- Now it remains only to show that in each of the strips, except possibly the bottommost Type 2 strip, $|\text{number of men} - \text{number of women}| \leq \frac{n}{25\sqrt{T}}$.

We will follow the proof of Theorem 3.15, though the proof will be simplified by the fact that we only need to consider $\sqrt{T}$ many time steps.

We divide each strip into thin diagonals of width 1. Let the diagonal include the bottom but not the top boundary. Notice that for each value, a diagonal contains at most one grid point.

As in Theorem 3.15, we introduce the following notation w.r.t. diagonal d at time step τ, where we are conditioning on the outcome of step $\tau - 1$.

$$
\begin{aligned}
I(d,\tau) &= E[(\text{number of men at time } \tau - \text{number of women at time } \tau)] \\
X(d,\tau) &= (\text{number of men matching at time } \tau - \text{number of women matching at time } \tau) \\
&\quad - E[(\text{number of men matching at time } \tau - \text{number of women matching at time } \tau)] \\
Y(d,\tau) &= \text{number of men entering at time } \tau - \text{number of women entering at time } \tau \\
A(d,\tau) &= (\text{number of men matching at time } \tau + \text{number of women matching at time } \tau)/2.
\end{aligned}
$$

$I(d,\tau)$ is measured after the entry of the new agents at time τ but prior to the match for this step. Also, note that $Y(d,\tau) = 0$ if d is in a Type 2 strip.

In addition, observe that the imbalance $\text{Imb}(s)$ at the start of step t equals $\sum_{d\in s} I(d,t)$.

We observe that a match between two agents in distinct diagonals of the same strip will increment the (number of men $-$ number of women) in one diagonal and decrement it in the other. Thus there is a zero net change over all the diagonals in the strip due to the matches. However, as the agents all age by 1 unit during a step, some agents enter the strip and some leave, which can cause changes to the imbalance within a strip. However, the entry of new agents can introduce new imbalances. We will need to understand more precisely how these imbalances evolve.

It is convenient to number the diagonals as $d_1, d_2, d_3, \ldots$, in right to left order.

We recall the following claims from the proof of Theorem 3.15.

Claim A.1.1. *Let d_i and d_j be two diagonals in the same strip s. For brevity, let $I_i \triangleq I(d_i, \tau-1)$, $I_j \triangleq I(d_j, \tau-1)$, $A_i \triangleq A(d_i, \tau-1)$, $A_j \triangleq A(d_j, \tau-1)$, $X_i \triangleq X(d_i, \tau-1)$, $X_j \triangleq X(d_j, \tau-1)$. Finally, let R denote the maximum of the total number of men and the total number of women in the system at time $\tau-1$. Then the new imbalance on diagonal d_i, prior to every unmatched agent adding 1 to their age (which causes the agents on d_i to move to d_{i+1}), denoted by $I'(d_i, \tau)$, is given by:*

$$I'(d_i, \tau) =$$
$$I_i + X_i - \sum_{d_j \in s} \left[X_i \frac{(2A_j - I_j - X_j)}{4R} - X_j \frac{(2A_i - I_i - X_i)}{4R} + I_i \frac{(2A_j - I_j - X_j)}{4R} - I_j \frac{(2A_i - I_i - X_i)}{4R} \right];$$

and $I(d_i, \tau) = I'(d_{i-1}, \tau-1) + Y(d, \tau)$.

$X(d, \tau)$ and $Y(d, \tau)$ are generated at diagonal d at time τ and, by Claim A.1.1, at any subsequent time step, $X(d, \tau)$ and $Y(d, \tau)$ will be redistributed over other diagonals.

1. Due to the expected matching at time $\tau' \geqslant \tau$, each $X(d, \tau)$ and $Y(d, \tau)$ spreads to other diagonals in the same strip.

2. At the end of time step τ' the portions of $X(d, \tau)$ and $Y(d, \tau)$ present on diagonal d_i move to diagonal d_{i+1}.

Building on these observations, we will show our bound on the imbalance by bounding the total contribution from $X(\cdot, \tau)$ and $Y(\cdot, \tau)$ to strip s at time τ'.

Notice that $\sum_{d_i \in s} I'(d_i, \tau) = \sum_{d_i \in s} I(d_i, \tau - 1)$, for the coefficients multiplying X_i cancel, as they also do for I_i. Thus we can think of this process as redistributing the imbalance, but not changing the total imbalance.

Over time an imbalance $X(d_i, \tau)$ will be redistributed over many diagonals.

We write $X(d_i, \tau, d_j, \tau')$ to denote the portion of $X(d_i, \tau)$ on diagonal d_j at time τ'. d_j need not be in the same strip as d_i. Note that $\sum_{d_j} X(d_i, \tau, d_j, \tau') = X(d_i, \tau)$ for all $\tau' \geqslant \tau$. $Y(d_i, \tau, d_j, \tau')$ is defined analogously.

For the purposes of the following claim, we treat the final strip as a single diagonal, and in addition ignore the fact that people depart at age T (which means that once an imbalance appears in this strip it remains there). The reason this strip is different is that it covers the whole of the bottom boundary and so is the only strip from which people leave the system by aging out.

Claim A.1.2. *For all ℓ, for all $i < k$, and for all $\tau' \geqslant \tau$, $\left|\sum_{j>\ell} X(d_i, \tau, d_j, \tau')\right| \leqslant \left|\sum_{j>\ell} X(d_k, \tau, d_j, \tau')\right|$. The same property holds for the $Y(d_i, \tau, d_j, \tau')$.*

Later, we will show a common bound B on the sums $\left|\sum_{i \leqslant j \leqslant k} X(d_j, \tau)\right|$, which holds for all d_i and d_k in the same strip and all τ.

With this bound and Claim A.1.2 in hand, for each strip s, we can bound the contribution of the $X(d_i, \tau, d_j, \tau')$ summed over all d_i and over $d_j \in s$ by $2B$.

Claim A.1.3. *For all $\tau' \geqslant \tau$, for every strip s, $\left|\sum_{d_i; d_j \in s} X(d_i, \tau, d_j, \tau')\right| \leqslant 2B$.*

This allows us to obtain the bound the imbalance in a strip s at any time $\tau' \leq \sqrt{T}$ by considering the contributions of $\left|\sum_{d_i; d_j \in s} X(d_i, \tau, d_j, \tau')\right|$ and $\left|\sum_{d_i; d_j \in s} Y(d_i, \tau, d_j, \tau') \cdot Y(d_i, \tau)\right|$ at all possible previous times (which is at most $\sqrt{T}$ time).

Regarding the contribution of Y, we also have the following claim,

Claim A.1.4. *With probability at least $1 - \frac{1}{n^{2c+1}}$, for all $\tau' \geqslant \tau$, for every strip s,*

$\left|\sum_{d_i; d_j \in s} Y(d_i, \tau, d_j, \tau')\right| \leqslant 2\sqrt{\frac{3n}{2} \ln\left(2Tn^{3c+1}\right)}$.

Thus,

$$|\operatorname{Imb}(s,\tau')| \leq \left[2B + 2\sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)}\right]\sqrt{T} \tag{A.14}$$

We now calculate B.

Claim A.1.5. *For any diagonal d_j and any d_i and d_k that lie in the same strip, at any time $\tau \leq \sqrt{T}$,*

$$\Pr\left[\left|\sum_{i\leq j\leq k} X(d_j,\tau)\right| \geq 2\sqrt{\frac{n(1+\epsilon)^2\sqrt{3}\ln\left(16T\sqrt{T}(\sqrt{T}+1)n^{c+1}\right)}{0.48\sqrt{T}}}\right] \leq \frac{1}{4n^{c+1}}.$$

Proof. The agents enter with one of T values chosen uniformly at random and are equally likely to be men or women. Hence, for all $\tau \leq \sqrt{T}$ time steps, for each value v,

$$\Pr\left[\text{At most } \frac{n(1+\epsilon)}{2T} \text{ men enter with value } v\right] \geq 1 - \tau T e^{-\frac{\epsilon^2 n}{6T}} \geq 1 - T\sqrt{T}e^{-\frac{\epsilon^2 n}{6T}}.$$

Call this event $\mathcal{E}_m$. Similarly,

$$\Pr\left[\text{At most } \frac{n(1+\epsilon)}{2T} \text{ men enter with value } v\right] \geq 1 - \tau T e^{-\frac{\epsilon^2 n}{6T}} \geq 1 - T\sqrt{T}e^{-\frac{\epsilon^2 n}{6T}}.$$

Call this event $\mathcal{E}_w$. Henceforth we condition on $\mathcal{E}_m$ and $\mathcal{E}_w$.

Consider some time τ. At this time, let $M = \max\{\text{total number of men, total number of women}\}$ while m = number of men in strip s and w = number of men in strip s. Using Lemmas 3.4 and 3.3, we obtain the following bound on the deviation from the expected number of the number of men in s matched in a given time step, τ:

$$\Pr\left[\left|\begin{array}{c}\text{number of men matched}\\ -E[\text{number of men matched}]\end{array}\right| > \frac{mw\delta}{M}\right] \leq 2e^{-mw\delta^2/3M}. \tag{A.15}$$

We will later prove the following claim,

Claim A.1.6. *For all time* $0 \le t \le \sqrt{T}$, $0.12nt \le M \le nt$ *for all* $t \le \sqrt{T}$, *with probability at least* $1 - \frac{1}{5n^{2c+1}}$.

Call this event $\mathcal{E}_M$. Henceforth, we further condition on $\mathcal{E}_M$.

Thus, from equation A.15, we obtain

$$\Pr\left[\left| \begin{array}{c} \text{number of men matched} \\ -E[\text{number of men matched}] \end{array} \right| > \frac{mw\delta}{0.12nt}\right] \le 2e^{-mw\delta^2/3nt}.$$

Setting $\delta = \left[\frac{3nt}{mw}\ln(n^c A(n,T))\right]^{1/2}$, we obtain

$$\Pr\left[\left| \begin{array}{c} \text{number of men matched} \\ -E[\text{number of men matched}] \end{array} \right| > \sqrt{\frac{mw\sqrt{3}\ln(n^c A(n,T))}{0.12nt}}\right] \le \frac{2}{A(n,T)n^c}.$$

We will specify $A(n,T)$ later. It is easy to see that because of $\mathcal{E}_m$ and $\mathcal{E}_w$, $m \le nt(1+\epsilon)/2\sqrt{T}$ and $w \le nt(1+\epsilon)/2\sqrt{T}$. So we obtain,

$$\Pr\left[\left| \begin{array}{c} \text{number of men matched} \\ -E[\text{number of men matched}] \end{array} \right| > \sqrt{\frac{nt(1+\epsilon)^2\sqrt{3}\ln(n^c A(n,T))}{0.48T}}\right] \le \frac{2}{A(n,T)n^c}.$$

Since $t \le \sqrt{T}$,

$$\Pr\left[\left| \begin{array}{c} \text{number of men matched} \\ -E[\text{number of men matched}] \end{array} \right| > \sqrt{\frac{n(1+\epsilon)^2\sqrt{3}\ln(n^c A(n,T))}{0.48\sqrt{T}}}\right] \le \frac{2}{A(n,T)n^c}.$$

We can perform the same argument for the women to obtain,

$$\Pr\left[\left|\begin{array}{c}\text{number of women matched}\\ -E[\text{number of women matched}]\end{array}\right| > \sqrt{\frac{n(1+\epsilon)^2\sqrt{3}\ln(n^c A(n,T))}{0.48\sqrt{T}}}\right] \le \frac{2}{A(n,T)n^c}.$$

From this it immediately follows that

$$\Pr\left[\left|\sum_{d\in S} X(d,\tau)\right| > 2\sqrt{\frac{n(1+\epsilon)^2\sqrt{3}\ln(n^c A(n,T))}{0.48\sqrt{T}}}\right] \le \frac{4}{A(n,T)n^c}.$$

where S is any subset of the diagonals in strip s. We have to consider $\sqrt{T}$ many possible times, T many diagonals d_j and up to $(\sqrt{T}+1)$ many strips. Thus setting $A(n,T) = 16T\sqrt{T}(\sqrt{T}+1)n$, proves the claim. □

From equation (A.14) we obtain, for every strip s and $\tau \le \sqrt{T}$, the following bound on $|\operatorname{Imb}(s,\tau')|$:

$$\begin{aligned}|\operatorname{Imb}(s,\tau')| &\le \left[2B + 2\sqrt{\frac{3n}{2}\ln\left(2Tn^{3c+1}\right)}\right]\sqrt{T}\\ &\le 8\sqrt{n(1+\epsilon)^2\sqrt{T}\ln\left(16T\sqrt{T}(\sqrt{T}+1)n^{c+1}\right)} + \sqrt{6nT\ln\left(2Tn^{3c+1}\right)}.\end{aligned}$$

We are conditioning on $\mathcal{E}_m$, $\mathcal{E}_w$ and $\mathcal{E}_M$. Set $\epsilon = \left[\frac{6T}{n}\ln\left(4T\sqrt{T}n^{3c+1}\right)\right]^{1/2}$. We choose constraints so that $\epsilon \le 1$. So, for every strip s and $\tau \le \sqrt{T}$,

$$\begin{aligned}&\Pr\left[\left|\operatorname{Imb}(s,\tau')\right| \le 16\sqrt{n\sqrt{T}\ln\left(16T\sqrt{T}(\sqrt{T}+1)n^{c+1}\right)} + \sqrt{6nT\ln\left(2Tn^{3c+1}\right)}\right]\\ &\ge \left(1 - \frac{1}{4n^{c+1}} - \frac{1}{n^{2c+1}}\right)\left(1 - \frac{1}{5n^{2c+1}} - \frac{1}{2n^{2c+1}}\right).\end{aligned}$$

Then we obtain,

$$\Pr\left[\left|\mathrm{Imb}(s,\tau')\right| \le 16\sqrt{n\sqrt{T}\ln\Big(16n^{c+1}T\sqrt{T}(\sqrt{T}+1)\Big)} + \sqrt{6nT\ln\left(2Tn^{3c+1}\right)}\right]$$
$$\ge \left(1-\frac{1}{4n^{c+1}}-\frac{1}{n^{2c+1}}\right)\left(1-\frac{7}{10n^{2c+1}}\right)] \ge 1-\frac{1}{n^{c+1}}.$$

We desire that

$$16\sqrt{n\sqrt{T}\ln\Big(16n^{c+1}T\sqrt{T}(\sqrt{T}+1)\Big)} + \sqrt{6nT\ln\left(2Tn^{3c+1}\right)} \le \frac{n}{25\sqrt{T}}$$

for which it suffices (by constraints (3.4)) that

$$32\cdot 25\cdot T\sqrt{(3c+3)}\ln n \le \sqrt{n}.$$

or,

$$n \ge (3c+3)(25\cdot 32\cdot T\ln n)^2,$$

which is also provided by the constraints (3.4). □

It now remains to prove Claim A.1.6. We proceed exactly as in the proof of Theorem 3.9.

Proof. (of Claim A.1.6).

Let's consider those agents that enter at times in the range $[0,t]$ for some $t \leqslant \sqrt{T}$. We want to lower bound the number of these agents who are present in the pool for the match at time t.

Henceforth, we will only consider men with values in the range $[T+\sqrt{T}, 2T)$. Among these men, consider those who have been in the pool for t' time, where $0 \leqslant t' < \sqrt{T}$. Let $p_i^{t'}$ be the probability that during their t'-th time step, the men in strip i are offered a match in their own strip. Even if all these men were still present in the matching pool,

$$\Pr\left[\#\text{ of these men matched in strip } i \text{ at age } t' \leqslant \frac{n(1+\delta)(1+\epsilon)\cdot p_i\cdot w_i}{2T}\right] \ge 1-e^{-\frac{\delta^2 np_i w_i}{6T}},$$

where w_i is the horizontal width of strip i occupied by these men when aged t'. For every Type 1 strip, $w_i \leqslant \sqrt{T}$. For the one Type 2 strip, since all values are at least $T + \sqrt{T}$, for ages up to $\sqrt{T}$, $w_i \leqslant \sqrt{T}$. By applying $\overline{\mu} = \frac{n(1+\epsilon)\max\{p_i, \frac{1}{T}\}\sqrt{T}}{2T}$ in Lemma 3.5, it follows that:

$$\Pr\left[\#\text{ of these men matched in strip } i \text{ at age } t' \leqslant \frac{n(1+\delta)(1+\epsilon)\cdot\max\{p_i, \frac{1}{T}\}}{2\sqrt{T}}\right]$$
$$\geq 1 - e^{-\frac{\delta^2 n \max\{p_i, \frac{1}{T}\}\sqrt{T}}{6T}} \geq 1 - e^{-\frac{\delta^2 n}{6T^{1.5}}},$$

The sum of the match probabilities—the p_i's— is at most 1. Notice that at any fixed time we only need to consider $\sqrt{T}$ strips, because at any time step, the men we are considering will occupy only $\sqrt{T}$ many strips. This implies $\sum \max\{p_i, \frac{1}{T}\} \leq 1 + \frac{1}{\sqrt{T}}$. Therefore,

$$\Pr\left[\begin{array}{l}\#\text{ of these men being matched}\\ \text{over all the strips at age } t'\end{array} \leqslant \frac{(1+\frac{1}{\sqrt{T}})n(1+\delta)(1+\epsilon)}{2\sqrt{T}}\right] \geq 1 - \sqrt{T}\cdot e^{-\frac{\delta^2 n}{6T^{1.5}}}.$$

Hence, we can bound the probability of the number of men who entered at time $t - \Delta + 1$ and left by time t , for any $\Delta \leqslant t$, as follows:

$$\Pr\left[\begin{array}{l}\#\text{ men being matched}\\ \text{in their first } \Delta \text{ steps}\end{array} \leqslant \frac{(1+\frac{1}{\sqrt{T}})n\Delta(1+\delta)(1+\epsilon)}{2\sqrt{T}}\right] \geq 1 - \Delta\sqrt{T}\cdot e^{-\frac{\delta^2 n}{6T^{1.5}}}.$$

Consequently, we can bound the probability for the number of men that enter in the time interval $[0, t-1]$ and are matched no later than time $t-1$ as follows:

$$\Pr\left[\begin{array}{l}\#\text{ men who entered and were}\\ \text{matched in a } t-1 \text{ time window}\end{array} \leqslant \frac{(1+\frac{1}{\sqrt{T}})n(1+\delta)(1+\epsilon)(t-1)}{4}\right] \geq 1 - \frac{1}{2}t^2\sqrt{T}e^{-\frac{\delta^2 n}{6T^{1.5}}}.$$

We set $\delta = \left[\frac{6T^{1.5}}{n} \ln\left(10n^{2c+1}T^{1.5}\right)\right]^{1/2}$. Note that $t \leq \sqrt{T}$.

We already chose $\epsilon = \left[\frac{6T}{n} \ln\left(4T\sqrt{T}n^{2c+1}\right)\right]^{1/2}$. By constraint (3.4), $c \geq 1$, $400 \leq T \leq n$, and $n \geq 96T^2(2c+3)\ln n$, $\delta \leq 1/4$ and $\epsilon \leq 1/64$. This yields the bound:

$$\Pr\left[\begin{array}{c}\text{\# men who entered in a } t-1 \\ \text{window being matched}\end{array} \geqslant \frac{\frac{65}{64}\cdot\frac{5}{4}nt}{4}\right] \geq 1 - \frac{1}{20n^{2c+1}}.$$

Since we have been conditioning on $\mathcal{E}$, this bound holds with probability at least $1 - \frac{1}{10n^{2c+1}}$. The same bound applies to the women.

Recalling that we excluded the men with values less than $T + \sqrt{T}$, this yields the following lower bound on the total population size, at time t:

$$nt - \frac{n(1+\epsilon)}{\sqrt{T}} - 0.635nt \geq 0.25nt,$$

with probability at least $1 - \frac{1}{5n^{2c+1}}$.

Thus $0.12nt \leq M \leq nt$ for all $t \leq T$, with probability at least $1 - \frac{1}{5n^{2c+1}}$, which proves the claim. $\square$

A.2 Further Detailed Data from Numerical Simulations in Section 3.7

Here we provide the average population size and average loss we obtained from every run of our numerical simulations.[4].

A.2.1 Reasonable Strategy

- Discrete Setting:

  $n = 500, T = 100$:

Run number	Average population size	Average Loss/T
1	1544.1	10.07
2	1507.7	9.96
3	1558.7	10.13
4	1458.1	9.8
5	1456.4	9.82
6	1565.5	10.16
7	1487.6	9.91
8	1567	10.17
9	1499.5	9.93
10	1508.3	9.99

Table A.1: Reasonable Strategy ($n = 500$, $T = 100$): Average population size and loss

Average Population Size = 1511.7 ± 3.7%.

Average Loss/T = 9.99 ± 1.9%.

[4] every run was for 2000 iterations

- Continuum Setting (= 500, T = 100):

  Average Population Size = 1180.9.

  Average Loss/T = 8.89.

A.2.2 Modified Reasonable Strategy

- Discrete Setting:

  - $n = 500$, $T = 100$:

Run number	Average population size	Average Loss/T
1	4733.1	22.04
2	4735	22.06
3	4666.8	21.82
4	4827.4	22.38
5	4726	22.04
6	4685.4	21.87
7	4785.1	22.25
8	4732.8	22.06
9	4681.4	21.89
10	4721.2	22.02

Table A.2: Modified Reasonable Strategy ($n = 500$, $T = 100$): Average population size and loss

Average Population Size = 4747.1 ± 1.7%.

Average Loss/T = 22.1 ± 1.7%.

  - $n = 500$, $T = 200$:

Run number	Average population size	Average Loss/T
1	6972.6	32.39
2	7194.9	33.17
3	7210.3	33.19
4	7148.8	33.01
5	7336.1	33.67
6	7085.7	32.79
7	7153.2	32.97
8	7119.2	32.94
9	7195.2	33.11
10	7086.6	32.79

Table A.3: Modified Reasonable Strategy ($n = 500$, $T = 200$): Average population size and loss

Average Population Size = 7154.4 ± 2.6%.

Average Loss/T = 33.03 ± 2%.

– $n = 500$, $T = 300$:

Run number	Average population size	Average Loss/T
1	8963.4	41.25
2	8930.5	41.21
3	8937	41.17
4	8744.5	40.49
5	8889.2	41.07
6	8798.6	40.75
7	8854.8	40.89
8	8783.7	40.68
9	8682.9	40.34
10	8705.3	40.37

Table A.4: Modified Reasonable Strategy ($n = 500$, $T = 300$): Average population size and loss

Average Population Size = 8823.2 ± 1.6%.

Average Loss/T = 40.8 ± 1.2%.

- $n = 500$, $T = 400$:

Run number	Average population size	Average Loss/T
1	10554.7	48.4
2	10281	47.22
3	10060.9	46.52
4	10168.1	47
5	10748.9	48.9
6	10295.3	47.39
7	10185.9	46.86
8	10094.6	46.75
9	10555.1	48.3
10	10185.3	46.99

Table A.5: Modified Reasonable Strategy ($n = 500$, $T = 400$): Average population size and loss

Average Population Size = 10404.9 ± 3.4%.

Average Loss/T = 47.71 ± 2.5%.

– $n = 500$, $T = 500$:

Run number	Average population size	Average Loss/T
1	11186.2	52.09
2	11268.1	52.38
3	11530	53.44
4	12209.8	55.8
5	12198.1	55.63
6	12027.3	55.07
7	11710.77	54.03
8	11372.4	52.78
9	11493.2	53.27
10	11378.5	52.77

Table A.6: Modified Reasonable Strategy ($n = 500$, $T = 500$): Average population size and loss

Average Population Size = 11698 ± 4.4%.

Average Loss/T = 53.95 ± 3.5%.

In summary:

n	T	Average population size	Average Loss/T
500	100	4747.1 ± 1.7%	22.1 ± 1.7%
500	200	7154.4 ± 2.6%	33.03 ± 2%
500	300	8823.2 ± 1.6%	40.8 ± 1.2%
500	400	10404.9 ± 3.4%.	47.71 ± 2.5%
500	500	11698 ± 4.4%	53.95 ± 3.5%

Table A.7: Modified Reasonable Strategy (summary): Average population size and loss

- Continuum Setting ($n = 500$, $T = 100$):

Average Population Size = 4484.8

Average Loss/T = 20.85.

Bibliography

Abadi, H. K. and Prabhakar, B. (2017). Stable matchings in metric spaces: Modeling real-world preferences using proximity. *CoRR*, abs/1710.05262.

Abdulkadiroglu, A. and Sönmez, T. (2013). Matching markets: Theory and practice. *Advances in Economics and Econometrics*, 1:3–47.

Abdulkadiroğlu, A., Pathak, P. A., and Roth, A. E. (2005). The new york city high school match. *American Economic Review*, 95(2):364–367.

Adachi, H. (2003). A search model of two-sided matching under nontransferable utility. *Journal of Economic Theory*, 113(2):182–198.

Akbarpour, M., Li, S., and Gharan, S. O. (2020). Thickness and information in dynamic matching markets. *Journal of Political Economy*, 128(3).

Amanatidis, G., Birmpas, G., Christodoulou, G., and Markakis, E. (2017). Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 545–562.

Ashlagi, I., Braverman, M., Kanoria, Y., and Shi, P. (2019). Clearing matching markets efficiently: Informative signals and match recommendations. *Management Science*, 66(5):2163–2193.

Ashlagi, I., Kanoria, Y., and Leshno, J. D. (2017). Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy*, 125(1).

Avis, D. (1983). A survey of heuristics for the weighted matching problem. *Networks*, 13(4):475–493.

Babaioff, M., Lavi, R., and Pavlov, E. (2006). Single-value combinatorial auctions and implementation in undominated strategies. In *SODA*, volume 6, pages 1054–1063. Citeseer.

Baccara, M., Lee, S., and Yariv, L. (2020). Optimal dynamic matching. *Theoretical Economics*, 15:1221––1278.

Bearden, J. (2006). A new secretary problem with rank-based selection and cardinal payoffs. *J. of Math.l Psych.*, 50:58–9.

Biggs, N. (1988). Matching theory (annals of discrete mathematics 29).

Biró, P. (2017). Applications of matching models under preferences.

Biró, P. and Klijn, F. (2013). Matching with couples: a multidisciplinary survey. *International Game Theory Review*, 15(02):1340008.

Bloch, F. and Ryder, H. (2000). Two-sided search, marriages, and matchmakers. *International Economic Review*, 41(1):93–116.

Blum, A., Hajiaghayi, M., Ligett, K., and Roth, A. (2008). Regret minimization and the price of total anarchy. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 373–382.

Blum, A. and Mansour, Y. (2007). Learning, regret minimization, and equilibria. *Algorithmic Game Theory*.

Burdett, K. and Coles, M. G. (1997). Marriage and Class*. *The Quarterly Journal of Economics*, 112(1):141–168.

Burdett, K. and Coles, M. G. (1999). Long-term partnership formation: Marriage and employment. *The Economic Journal*, 109(456):307–334.

Burdett, K. and Wright, R. (1998). Two-sided search with nontransferable utility. *Review of Economic Dynamics*, 1(1):220–245.

Chade, H., Eeckhout, J., and Smith, L. (2017). Sorting through search and matching models in economics. *Journal of Economic Literature*, 55(2):493–544.

Chau, S. C.-K., Shen, S., and Zhou, Y. (2020). Decentralized ride-sharing and vehicle-pooling based on fair cost-sharing mechanisms. *IEEE Transactions on Intelligent Transportation Systems*, 23(3):1936–1946.

Coles, P., Kushnir, A., and Niederle, M. (2013). Preference signaling in matching markets. *American Economic Journal: Microeconomics*, 5(2):99–134.

Damiano, E., Li, H., and Suen, W. (2005). Unravelling of dynamic sorting. *The Review of Economic Studies*, 72(4):1057–1076.

Diebold, F., Aziz, H., Bichler, M., Matthes, F., and Schneider, A. (2014). Course allocation via stable matching. *Business & Information Systems Engineering*, 6:97–110.

Diebold, F. and 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.

Dobzinski, S., Ron, S., and Vondrák, J. (2022). On the hardness of dominant strategy mechanism design. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 690–703.

Dubins, L. E. and Freedman, D. A. (1981). Machiavelli and the gale-shapley algorithm. *The American Mathematical Monthly*, 88(7):485–494.

Eeckhout, J. (1999). Bilateral search and vertical heterogeneity. *International Economic Review*, 40(4):869–887.

Fabrikant, A. and Papadimitriou, C. H. (2008). The complexity of game dynamics: Bgp oscillations, sink equilibria, and beyond. In *SODA*, volume 8, pages 844–853. Citeseer.

Feigenbaum, J., Krishnamurthy, A., Sami, R., and Shenker, S. (2003). Hardness results for multicast cost sharing. *Theoretical Computer Science*, 304(1-3):215–236.

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, 4(3):282–9.

Finocchiaro, J., Maio, R., Monachou, F., Patro, G. K., Raghavan, M., Stoica, A.-A., and Tsirtsis, S. (2021). Bridging machine learning and mechanism design towards algorithmic fairness.

Friedman, E. and Shenker, S. (1998). Learning and implementation on the internet. Technical report, Working paper.

Gale, D. et al. (2001). The two-sided matching problem. origin, development and current issues. *International Game Theory Review*, 3(2/3):237–252.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.

Gimbert, H., Mathieu, C., and Mauras, S. (2019). Incentives in popularity-based random matching markets.

Gonczarowski, Y. A., Nisan, N., Kovalio, L., and Romm, A. (2019). Matching for the israeli "mechinot" gap-year programs: Handling rich diversity requirements. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 321, New York, NY, USA. Association for Computing Machinery.

Gonczarowski, Y. A., Nisan, N., Ostrovsky, R., and Rosenbaum, W. (2015). A stable marriage requires communication. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, page 1003–1017, USA. Society for Industrial and Applied Mathematics.

Gusfield, D. and Irving, R. W. (1989). *The stable marriage problem: structure and algorithms*. MIT press.

Hakimov, R. and Kübler, D. (2019). Experiments on matching markets: A survey. Technical report, WZB Discussion Paper.

Hassidim, A., Romm, A., and Shorrer, R. I. (2017). Redesigning the israeli psychology master's match. *American Economic Review*, 107(5):205–09.

Immorlica, N. and Mahdian, M. (2015). Incentives in large random two-sided markets. *ACM Trans. Econ. Comput.*, 3(3).

Iwama, K. and Miyazaki, S. (2008). A survey of the stable marriage problem and its variants. In *International Conference on Informatics Education and Research for Knowledge-Circulating Society (icks 2008)*, pages 131–136.

Kanoria, Y., Min, S., and Qian, P. (2021). In which matching markets does the short side enjoy an advantage?

Karp, R. M., Vazirani, U. V., and Vazirani, V. V. (1990). An optimal algorithm for on-line bipartite matching. In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing*, STOC '90, page 352–358, New York, NY, USA. Association for Computing Machinery.

Kleinberg, J. (2000a). The small-world phenomenon: An algorithmic perspective. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 163–170.

Kleinberg, J. (2001). Small-world phenomena and the dynamics of information. *Advances in neural information processing systems*, 14.

Kleinberg, J. (2004). The small-world phenomenon and decentralized search. *SiAM News*, 37(3):1–2.

Kleinberg, J. (2006). Complex networks and decentralized search algorithms. In *Proceedings of the International Congress of Mathematicians (ICM)*, volume 3, pages 1019–1044. Citeseer.

Kleinberg, J. M. (2000b). Navigation in a small world. *Nature*, 406(6798):845–845.

Knuth, D. E. (1976). *Mariages stables et leurs relations avec d'autres problèmes combinatoires : introduction à l'analyse mathémathique des algorithmes*. Les Presses de l'Université de Montréal.

Knuth, D. E. (1996). *Stable Marriage and Its Relation to Other Combinatorial Problems: An Introduction to the Mathematical Analysis of Algorithms*, volume 10. CRM Proceedings & Lecture Notes.

Knuth, D. E., Motwani, R., and Pittel, B. (1990). Stable husbands. *Random Structures & Algorithms*, 1(1):1–14.

Kojima, F. and Pathak, P. A. (2009). Incentives and stability in large two-sided matching markets. *American Economic Review*, 99(3):608–27.

Kupfer, R. (2020). The influence of one strategic agent on the core of stable matchings. In *WINE*.

Lauermann, S. and Nöldeke, G. (2014). Stable marriages and search frictions. *Journal of Economic Theory*, 151:163–195.

Lee, S. (2016). Incentive compatibility of large centralized matching markets. *The Review of Economic Studies*, 84(1):444–463.

Manlove, D. (2013). *Algorithmics of matching under preferences*, volume 2. World Scientific.

McNamara, J. M. and Collins, E. (1990). The job search problem as an employer-candidate game. *Journal of Applied Probability*, pages 815–827.

Mehta, A. et al. (2013). Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science*, 8(4):265–368.

Mertens, S. (2005). Random stable matchings. *Journal of Statistical Mechanics: Theory and Experiment*, 2005.

Milgram, S. (1967). The small world problem. *Psychology today*, 2(1):60–67.

Nisan, N., Schapira, M., and Zohar, A. (2008). Best-reply mechanisms. Technical report, Citeseer.

nrmp.org (2021). Results and data, 2021 main residency match.

Panconesi, A. and Srinivasan, A. (1997). Randomized distributed edge coloring via an extension of the Chernoff-Hoeffding bounds. *SIAM J. Comput.*, 26(2):350–368.

Papadimitriou, C., Schapira, M., and Singer, Y. (2008). On the hardness of being truthful. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 250–259. IEEE.

Pittel, B. (1989). The average number of stable matchings. *SIAM Journal on Discrete Mathematics*, 2(4):530–549.

Pittel, B. (1992). On Likely Solutions of a Stable Marriage Problem. *The Annals of Applied Probability*, 2(2):358 – 401.

Pittel, B. (2019). On likely solutions of the stable matching problem with unequal numbers of men and women. *Mathematics of Operations Research*, 44(1):122–146.

Pittel, B., Shepp, L., and Veklerov, E. (2008). On the number of fixed pairs in a random instance of the stable marriage problem. *SIAM Journal on Discrete Mathematics*, 21(4):947–958.

Ren, J., Xia, F., Chen, X., Liu, J., Hou, M., Shehzad, A., Sultanova, N., and Kong, X. (2021). Matching algorithms: Fundamentals, applications and challenges. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 5(3):332–350.

Rios, I., Larroucau, T., Parra, G., and Cominetti, R. (2021). Improving the chilean college admissions system. *Operations Research*, 69(4):1186–1205.

Rogerson, R., Shimer, R., and Wright, R. (2005). Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43:959––988.

Roth, A. E. (1982). The economics of matching: Stability and incentives. *Mathematics of operations research*, 7(4):617–628.

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. (2008). Deferred acceptance algorithms: History, theory, practice, and open questions. *international Journal of game Theory*, 36:537–569.

Roth, A. E. and 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. and Sotomayor, M. (1990). A study in game-theoretic modeling and analysis. *Econometric Society Monographs*, 18.

Roth, A. E. and Sotomayor, M. (1992). Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541.

Roughgarden, T. (2008). An algorithmic game theory primer.

Shimer, R. and Smith, L. (2000). Assortative matching and search. *Econometrica*, 68(2):343–369.

Shorrer, R. I. (2019). Simultaneous search: Beyond independent successes. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 347–348, New York, NY, USA. Association for Computing Machinery.

Smith, L. (2006). The marriage model with search frictions. *Journal of political Economy*, 114(6):1124–1144.

Suppes, P. (1958). Michael o. rabin. effective computability of winning strategies. contributions to the theory of games, volume iii, annals of mathematics studies number 39, princeton university press, princeton, new jersey, 1957, pp. 147–157. *The Journal of Symbolic Logic*, 23(2):224–224.