Stable Matchings with Switching Costs*

Boris Pittel† Kirill Rudov‡

June 27, 2025

Abstract

Traditional matching theory, and its canonical stability notion, assumes that agents can freely and costlessly switch partners. Without switching costs, large matching markets have a small number of stable matchings. In reality, however, switching partners is often costly: in labor markets, employees may need to move; in marriage markets, divorce frequently carries financial and emotional burdens. We study the impacts of switching costs and find that they can dramatically expand the set of stable matchings, even in large markets, and with vanishingly small costs. We precisely characterize the threshold of switching costs that triggers transformative expansion: an explosion in the number of stable matchings. From a market design perspective, accounting for switching costs, stability allows for significantly more room for policy interventions than previously thought. Our results provide insights into competitive forces in markets with imbalances between supply and demand.

Keywords: Matching Markets, Stability, Market Design.

*We thank Federico Echenique, Haluk Ergin, Michael Ostrovsky, and Leeat Yariv for helpful comments.

†Ohio State University; pittel.1@osu.edu

‡University of California, Berkeley; krudov@berkeley.edu

1 Introduction

1.1 Overview

Stability is a central notion for two-sided matching markets. It is defined by the absence of blocking pairs—pairs of market participants who prefer each other over their current partners (Gale and Shapley, 1962). The core idea is that some participants may profitably deviate and circumvent the unstable matching in place by forming blocking pairs. Implicit is the assumption that participants can leave their partners and switch to better ones at no cost. In practice, however, participants may encounter hurdles that preclude the formation of certain blocking pairs. Specifically, switching partners may entail non-trivial costs, arising from various sources: contractual obligations or reputational concerns in labor markets, legal and financial responsibilities or the presence of children in marriage, or psychological factors such as guilt, social pressures, and inertia, just to name a few. Despite the prevalence of switching costs, our understanding of their effects remains limited. In this paper, we study the impacts of switching costs on market outcomes. We introduce a new notion of stability accounting for such costs and analyze its implications. As we show, the presence of even small switching costs can dramatically alter predictions.

Our notion of stability with switching costs, which we refer to as ϵ-*stability*, is a natural extension of the traditional concept of stability. Namely, an agent who is matched would be willing to switch and match with another partner only if the benefits of the switch outweigh ϵ, the *switching cost*. We study the effects of switching costs in matching markets, using a random markets framework.[1] For balanced markets, with equal numbers of participants on both sides, we provide a sharp characterization of the threshold switching cost—as a function of the market size—beyond which the expected number of stable matchings expands dramatically (Theorem 1), a phenomenon we call the "explosion of the core." In particular, this threshold vanishes as the market size grows, suggesting that even minute switching costs may result in a significant increase of plausible stable outcomes. Notably, the same explosion, with the same vanishing threshold, holds uniformly across imbalanced markets (Theorem 2), providing a different angle on the now-conventional belief that even minimally imbalanced (Ashlagi, Kanoria, and Leshno, 2017), or large (Roth and Peranson, 1999), markets tend to have a small set of stable outcomes. Through computational experiments leveraging a link between our model and integer programming, we show that even in small imbalanced

[1]The significance of this framework for studying matching markets is highlighted by the 2023 SIGecom Test of Time Award won by Immorlica and Mahdian (2005) and Ashlagi, Kanoria, and Leshno (2017).

markets with small switching costs, multiplicity can be substantial, manifesting in many participants having significantly different partners across ϵ-stable matchings. We illustrate that our model is flexible to examine other aspects of interest, highlighting its role as a valuable benchmark.

Altogether, our results emphasize the relevance of switching costs in shaping market outcomes. Even small switching costs can significantly expand the set of plausible outcomes, introducing new forms of stability and fairness. This expansion suggests that switching costs may potentially serve as a useful tool in market design, enabling the targeting of outcomes that would otherwise be unattainable. Additionally, our findings uncover a possible challenge for empirical studies of decentralized markets that rely on the stability of realized outcomes, assuming a perfectly frictionless environment.

Model We consider a standard model of two-sided matching markets with cardinal utilities.[2] Cardinal utilities are needed in order to introduce switching costs. The model consists of n agents on the "short" side of the market, which we refer to as firms; and $n + k$ agents on the "long" side, referred to as workers. Here, the non-negative integer k quantifies the degree of market imbalance. Preferences of firms over workers, as well as those of workers over firms, are represented by utilities.

Our analysis centers on the concept of ϵ-stability. When there are no switching costs, $\epsilon = 0$, ϵ-stability corresponds to the standard notion of stability in matching theory. Our notion relates to concepts of approximate stability—these weakenings are commonly employed to establish the existence of "almost stable" outcomes when achieving exact stability is problematic.[3] On a more abstract level, ϵ-stability is reminiscent of ϵ-equilibria, in which no agent can improve by more than ϵ; see Nisan et al. (2007) and references therein.

Random Markets We consider a cardinal extension of the classical model of random matching markets. The classical model is generated by drawing a complete preference list for each firm, and each worker, independently and uniformly at random. Since its introduction (Knuth, 1976; Pittel, 1989), this framework has established itself as a fundamental

[2] A large body of literature examines matching settings in which participants' preferences are cardinal, including decentralized matching (e.g., Lauermann, 2013; Echenique, Robinson-Cortés, and Yariv, 2024), centralized matching (e.g., Lee, 2017; Fernandez, Rudov, and Yariv, 2022), assignment problems (e.g., Budish and Cantillon, 2012; Che and Tercieux, 2019), and others.

[3] For instance, see Abraham, Biró, and Manlove (2005), Che and Tercieux (2019), and Doğan and Ehlers (2021). Additionally, approximate stability has been studied from a computational perspective; see Caragiannis et al. (2021) and Chen, Skowron, and Sorge (2021), along with references there.

methodological tool for exploring matching markets; for a recent survey, see Leshno (2023).[4]

We model utility values using a natural value distribution that is widely applied in the literature. Specifically, the utility an agent receives when matched with an agent from the other side is drawn independently from an exponential distribution. In our model, the exponential distribution captures the intuitive idea that agents can more easily differentiate between their higher-ranked partners than those ranked lower on their list. In a later section, we demonstrate that our results translate to scenarios where agents find it less easy to distinguish between their higher-ranked partners, revealing additional patterns that emerge in such cases. Aside from our model, the exponential distribution has been extensively employed to represent individual valuations and utilities in many economic domains, such as auctions, finance, marketing, matching theory, and operations research, among others.[5] Moreover, the exponential distribution is often used to capture preferences in empirical settings; for instance, see Crawford and Yurukoglu (2012), Kim (2015), and Donna and Espín-Sánchez (2018).

Another notable advantage of working with the exponential distribution is that it renders our model tractable and amenable to analysis. Since we deal with ϵ-stability, tractability is a significant challenge. To the best of our knowledge, no general structural results exist for ϵ-stable matchings—it is known that stable matchings are ϵ-stable, but not much more. When there are no switching costs, the set of stable matchings has a very useful lattice structure, but there are no analogous results for ϵ-stable matchings. Yet, our approach offers a tractable framework for investigating their features, making it an ideal benchmark model.

Main Findings We show that switching costs have dramatic implications for matching markets. Switching costs translate into an abundance of stable matchings, even in large imbalanced markets, and even when these costs are very small. Our results contrast with the (by now) conventional view that even minimally imbalanced markets, or large markets, exhibit a small core; that is, they give rise to a very limited set of stable outcomes.

In our main results, we provide a sharp characterization of the critical threshold for

[4]More broadly, the random markets framework echoes approaches used in other economic settings. A notable example is voting, where random preference profiles are commonly employed to estimate the likelihood of paradoxes and evaluate voting rules; see Diss and Merlin (2021) and references therein. In game theory, there is a long tradition of using random profiles to study the properties of games and solution concepts; for instance, see McLennan (2005), Alon, Rudov, and Yariv (2024), and references cited within.

[5]For examples of studies in different domains, see Kim (2015), Donna and Espín-Sánchez (2018), and Hartline and Roughgarden (2014) for auctions; Chu, Leslie, and Sorensen (2011) and Zhou (2021) for finance and operations research; and Anshelevich, Das, and Naamad (2013), Lee and Yariv (2018), Kanoria, Saban, and Sethuraman (2018), Ashlagi et al. (2020), and Pęski (2022) for matching.

switching costs, beyond which the number of ϵ-stable matchings expands drastically. Consider first a sequence of balanced random markets, with n firms and n workers.[6] The behavior of the set of stable matchings depends on the magnitude of switching costs $\epsilon = \epsilon(n)$, as a function of market size n, where larger costs lead to larger cores. Theorem 1 delineates the transformative regime change for the core, through two separate cases. We show that, when costs vanish at least as fast as $\log n/n$, the expected number of ϵ-stable matchings can be sizable while still remaining within bounds similar to those when no costs are present (Proposition 1). In particular, this number can be larger than $e^{-1}n \log n$, the approximate expected number of stable matchings without switching costs (Pittel, 1989), but it remains polynomially large. Importantly, the behavior shifts dramatically when ϵ surpasses the $\log n/n$ threshold: the set of ϵ-stable matchings explodes. The number of ϵ-stable matchings surges faster than any polynomial function of n (Proposition 2).

Theorem 2 shows that switching costs crucially alter the predictions about core size in imbalanced markets. When switching costs are absent, even the slightest imbalance in large markets makes the core immediately collapse. In contrast, we demonstrate that even minimal switching costs result in a vast number of stable matchings. Specifically, consider markets with n firms and $n + k$ workers, where $k = k(n) \geq 1$ is an arbitrary market imbalance. We prove that the same critical threshold, $\log n/n$, applies *uniformly* across all possible imbalances $k = k(n)$. Thus, the critical switching cost threshold that triggers the expansion of the core remains consistent, regardless of market imbalance.

Notably, the explosion in the number of stable matchings occurs, even when switching costs are vanishingly small. Indeed, the critical threshold $\log n/n$ goes to zero as the market grows. Our results imply that, even in nearly frictionless imbalanced markets, a huge number of stable matchings may arise.

Figure 1 illustrates our findings, and shows that our large market results kick in at moderate market sizes. It demonstrates that the set of stable matchings expands significantly, even in relatively small markets with small switching costs. The market size of $n = 34$ in our simulations is due to computational constraints.[7] Switching costs of $\epsilon = 0.06$ and $\epsilon = 0.12$ are selected to lie below, and (respectively) slightly above, the reference value of $\log n/n \approx 0.10$. These costs are small compared to the average utility of 1—corresponding to the mean of

[6] For expositional simplicity, suppose that utilities are drawn from the standard exponential distribution. Our analysis is more general, and particularly applies to cases where agents encounter consistent "relative" switching costs.

[7] This analysis becomes computationally feasible because of a connection we establish to integer programming, discussed in more detail later. We employ 1,000 simulations for each market size.

the standard exponential distribution used for this exercise. Despite these small switching costs, the core is large across all imbalances, contrasting with the collapsed core in a no-cost scenario for imbalanced markets.

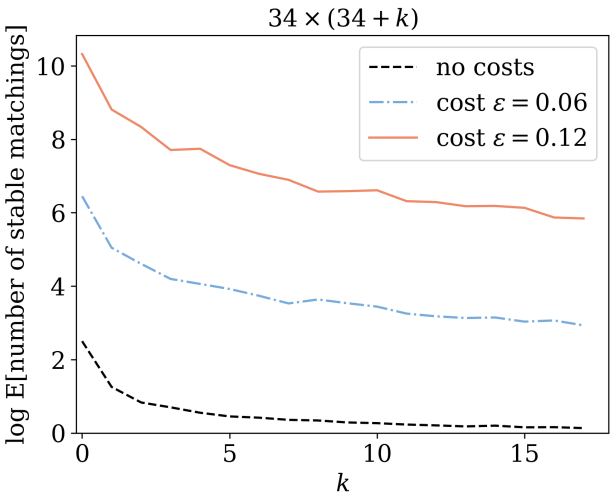


Figure 1: log [Number of ϵ-stable matchings]

Aside from illustrating our results for the number of stable matchings, the simulations serve to exhibit some additional patterns of interest in the presence of switching costs. We demonstrate that, even in small imbalanced markets with small switching costs, the expanded core involves a substantial fraction of agents experiencing significantly different outcomes across ϵ-stable matchings. Additionally, while these matchings still pair n firms with n workers, workers may be matched in one ϵ-stable matching but unmatched in another, and vice versa. This stands in contrast to the no-cost setting, where the Rural Hospital Theorem (McVitie and Wilson, 1970) states that the set of matched agents is the same in any stable matching, lending an air of inevitability to stability.

Our results have implications for the consequences of market imbalance, and the resulting effects of competitive pressure. Without switching costs, competition arising from even minimal imbalances yields an essentially unique stable matching (Ashlagi, Kanoria, and Leshno, 2017); for large imbalances, literally one stable matching remains (Pittel, 2019). This uniqueness leaves, effectively, no room for luck or disagreement among participants about who is matched with whom. With even small switching costs, the core might expand, bringing indeterminacy.

Implications Taken together, our results highlight the importance of accounting for switching costs, even in markets that appear frictionless. Even small switching costs can lead to a notable expansion of plausible matching outcomes, introducing indeterminacy into the process. Agents may experience significantly different outcomes across ϵ-stable matchings—or compared to a scenario without costs—including the phenomenon of being matched in some outcomes and unmatched in others. These discrepancies could potentially translate into distinct and diverging long-term outcomes for agents, similar to those documented in the labor economics literature on how initial job placements affect future career trajectories; for example, see Oyer (2006) and Kahn (2010).

Switching costs may serve as a valuable tool in market design. Abstracting away from implementation concerns, switching costs can offer a useful instrument for targeting certain outcomes for specific groups of market participants, influencing who gets matched with whom, determining who gets matched or remains unmatched, and addressing various distributional concerns (e.g., Combe et al., 2022), among others. In practice, the designed structure of switching costs may be more nuanced, applying differently to participants.[8]

Switching costs can shape perceptions of fairness in matching markets. Stability is often interpreted not only as a solution concept but also as a fairness notion.[9] In a particular sense, the notion of ϵ-stability delineates these two interpretations. When switching costs are small, ϵ-stability can still be seen as a fairness condition—market participants may have only a minimal amount, if any, of justified envy. Our results translate into the significant expansion of fair outcomes and introduce new perspectives on fairness. However, with larger switching costs, ϵ-stable matchings may no longer be considered approximately fair.[10]

Our results reveal a potential challenge for empirical studies on matching markets. A growing body of empirical literature uses the stability of realized matches as an identification assumption to infer preferences in decentralized settings, abstracting away from frictions; see

[8]Even in markets relying on a stable centralized clearinghouse, switching costs—imposed through restrictions or penalties—are often used to solidify the implemented stable matching, further discouraging post-match mobility and preventing unraveling. In the medical match, for example, trainees face various barriers to transferring, such as needing approval from their current program director. While these switching costs effectively encourage participants to stay, it may be worth exploring how this objective balances with other considerations. For instance, incentives for organizations to improve practices, participants' long-term satisfaction with their match or profession as preferences and market conditions evolve, and decisions to leave the market altogether (Wang et al., 2022).

[9]In the school choice context, stability corresponds to elimination of justified envy and is also regarded as a normative fairness criterion (Balinski and Sönmez, 1999; Abdulkadiroğlu and Sönmez, 2003).

[10]On a more speculative note, switching costs could serve as a quantitative measure of unfairness—given an outcome, one could assess fairness by examining a certain measure of the costs required to make this outcome stable. Larger measures would indicate less fair outcomes.

the surveys by Fox (2009), Chiappori and Salanié (2016), and Agarwal and Somaini (2023). However, our findings indicate that even small switching costs may cause realized matches to differ from those expected in a hypothetical frictionless environment. Although exact stability is an appealing structural assumption, additional attention to verifying the presence and impact of switching costs may benefit empirical work by helping to avoid possible biases in preference estimates.

Methodological Contributions We introduce a general framework to study the effects of switching costs in matching markets. Our framework allows us to derive integral formulas for key variables of interest, building on ideas that go back to Knuth (1976). These formulas are crucial for tractability and pave the way for precise analysis of ϵ-stable matchings, despite the absence of general results about their structure. To establish our main results, we develop techniques inspired by recent methods of Pittel (2019). In later sections, we briefly discuss additional questions that our model enables us to explore, which are part of our ongoing research. Furthermore, our framework is versatile and can be applied to study the impacts of switching costs in other matching settings, including one-sided markets.

Somewhat relatedly, there has been a quest, yet unresolved, to find a tractable model of random matching markets with ties. Most of the matching literature assumes that agents are never indifferent between any two partners. In practice, however, ties in preferences are widespread; see Erdil and Ergin (2017) and references therein. In a certain sense, our notion of ϵ-stability suggests that two partners are "incomparable" for an agent if their utilities differ by at most the switching cost ϵ. This incomparability induces a tolerance relation—a reflexive and symmetric binary relation—that violates transitivity, unlike indifference satisfying all three conditions. Thus, our random matching model can serve as a more nuanced version of one with ties, where ϵ captures the degree of incomparability among potential partners.

The integral-based approach is nevertheless not omnipotent. Technically, for standard stable matchings, certain nuanced questions call for deeper structural tools—such as rejection chains in the Deferred Acceptance algorithm (Gale and Shapley, 1962; McVitie and Wilson, 1971) or rotations (Irving and Leather, 1986)—which are grounded in the lattice structure unique to these matchings.

This raises the question: how do we even manage to calculate the number of ϵ-stable matchings to generate Figure 1? We show that ϵ-stable matchings can be obtained as the integer solutions of a specific system of linear inequalities—this system parallels that introduced by Vande Vate (1989) for traditional stability. For our simulations, we employ

integer programming to investigate the expanded set of stable matchings.

1.2 Related Literature

Our work belongs to the extensive literature on matching markets, which originates from the pioneering work of Gale and Shapley (1962) on stable matching. This literature has since evolved into a rich, multi-faceted, and interdisciplinary field, leading to many fundamental results and successful applications. For further details, see Roth and Sotomayor (1992) and Echenique, Immorlica, and Vazirani (2023).

We examine stability, in the context of switching costs, within a random markets framework. This framework stems from Knuth (1976) and Pittel (1989) and has become a leading methodological approach for understanding the key features of general matching markets (Leshno, 2023). Our paper relates to studies investigating the structure of stable matchings. An important recent finding is that even slight market imbalances not only confer a substantial advantage to agents on the short side but also cause the set of stable matchings to collapse, leading to an essentially unique stable matching (Ashlagi, Kanoria, and Leshno, 2017; Pittel, 2019).[11] We contribute by developing a benchmark model of markets with switching costs, introducing a relevant applied dimension and revealing novel patterns. In particular, we identify cost regimes under which nearly frictionless markets, with vanishingly small costs, exhibit numerous stable matchings.

Our stability notion resembles approximate solution concepts employed in various domains. More permissible stability notions are utilized in contexts where traditional stable outcomes do not exist or where exact stability is incompatible with other desiderata; see Abraham, Biró, and Manlove (2005), Che and Tercieux (2019), Doğan and Ehlers (2021), and Leshno (2022) for examples. Moreover, in computer science, approximate stability—conceptually similar to approximate equilibria in game theory (Nisan et al., 2007)—can provide computational advantages; see Caragiannis et al. (2021) and Chen, Skowron, and Sorge (2021), and discussion there. Shifting the focus to the practical implications of switching costs, we examine their effect on the set and structure of stable matchings.

[11] Motivated by the length restrictions on rank order lists in centralized markets (Roth and Peranson, 1999), Immorlica and Mahdian (2005) and Kojima and Pathak (2009) show that random markets with short, constant-sized, preference lists generate a small core, theoretically suggesting limited scope for manipulation in such markets. Nonetheless, this short list assumption generates many unmatched agents (Lee, 2017), see also Kanoria, Min, and Qian (2023). Additionally, there is growing evidence that reported preferences might differ significantly from truthful ones; see Hassidim et al. (2017), Echenique et al. (2022), Artemov, Che, and He (2023), and references therein. Several recent papers suggest that cores may be large in certain settings with specific preference structures (Biro et al., 2022; Hoffman, Levy, and Mossel, 2023; Rheingans-Yoo, 2024).

Starting with Vande Vate (1989), several papers have examined stable matchings in matching markets—having ordinal preferences by default—through the lens of linear programming; for instance, see Roth, Rothblum, and Vande Vate (1993) and Teo and Sethuraman (1998).[12] In a similar vein, we introduce a system of linear inequalities to analyze stable matchings in markets with cardinal preferences and switching costs.

Though less directly related, there is a recent and growing body of literature, both in economics and computer science, that explores various questions related to robustness in matching markets. For example, robust stable solutions in the presence of incomplete information about preferences (Aziz et al., 2020), stable outcomes tolerant to certain errors (Genc et al., 2019; Chen, Skowron, and Sorge, 2021; Gangam et al., 2022), and "locally stable" stable matchings (Rudov, 2024), to name just a few.

2 The Model

2.1 Basic Definitions

A two-sided *matching market* consists of a set of n firms $\mathcal{F} = \{f_i\}_{i\in[n]}$, where $[n] = \{1, 2, \ldots, n\}$; a set of $n + k$ workers $\mathcal{W} = \{w_j\}_{j\in[n+k]}$; and a profile of match utilities $\{u^f_{ij}, u^w_{ij}\}_{(i,j)\in[n]\times[n+k]}$.[13] For each pair (i, j), u^f_{ij} is firm f_i's utility from matching with worker w_j and u^w_{ij} is worker w_j's utility from matching with firm f_i. Throughout, we consider markets where all worker-firm pairs are mutually acceptable; that is, every agent prefers to be matched over remaining single. In addition, we focus on $k \geq 0$, unless stated otherwise, and refer to k as the *imbalance*.

A *matching* is a one-to-one function $\mu : \mathcal{F} \cup \mathcal{W} \to \mathcal{F} \cup \mathcal{W}$ that (i) assigns to each firm f_i either a worker or herself, $\mu(f_i) \in \mathcal{W} \cup \{f_i\}$; (ii) assigns to each worker w_j either a firm or himself, $\mu(w_j) \in \mathcal{F} \cup \{w_j\}$; (iii) is of order two, i.e, $\mu(f_i) = w_j$ if and only if $\mu(w_j) = f_i$. To simplify notation, we often denote $\mu(i) = j$ and $\mu(j) = i$ whenever $\mu(f_i) = w_j$.

Consider an arbitrary matching μ and some f_i and w_j that are not matched to one another at μ, i.e., $\mu(i) \neq j$. The pair (f_i, w_j) is said to form an *ϵ-blocking pair* if (i) f_i is single or *ϵ-prefers* w_j over her partner under μ, meaning that $u^f_{i,j} > u^f_{i,\mu(i)} + \epsilon$; and (ii) w_j is

[12] This line of work has also led to various stability notions for so-called fractional, or non-integral, matchings (Aziz and Klaus, 2019).

[13] For convenience, we label the two sides as firms and workers. In practice, they may correspond to doctors and hospitals, students and colleges, teachers and schools, men and women, actual firms and workers for labor markets with fixed wages (e.g., government and union jobs, see Hall and Krueger, 2012), or any other matching setting with no or limited transfers.

single or ϵ-prefers f_i, that is, $u^w_{i,j} > u^w_{\mu(j),j} + \epsilon$. A matching is *$\epsilon$-stable* if it has no ϵ-blocking pairs. For convenience, we often refer to ϵ—which is assumed to be non-negative—as the *switching cost*.

An ϵ-stable matching always exists. Specifically, a stable matching exists (Gale and Shapley, 1962), and any such matching is also ϵ-stable. Since single agents do not incur switching costs when forming an ϵ-blocking pair, all n firms are matched in any ϵ-stable matching.[14] Moreover, the concepts of stability and ϵ-stability coincide when $\epsilon = 0$.

Switching costs stem from a wide range of factors: financial, legal, emotional, or behavioral, such as status quo bias. For instance, in labor markets, even breaching an informal agreement may bring reputational damage; in relationships, ending a partnership can involve both emotional and bureaucratic challenges. Even small costs may deter market participants from pursuing certain otherwise beneficial deviations.

The following example illustrates the relevant concepts in our model.

Example 1. Consider a market with two firms and three workers. The following matrix provides a convenient representation of the agents' utilities:

$$\begin{array}{c} \\ f_1 \\ f_2 \end{array} \begin{array}{c} \begin{array}{ccc} w_1 & w_2 & w_3 \end{array} \\ \begin{pmatrix} 2,1 & 1,2 & 1.01,1 \\ 1.99,2 & 2,1 & 1,2 \end{pmatrix} \end{array}.$$

In this matrix, each row i corresponds to firm f_i, and each column j corresponds to worker w_j. Each (i,j)-entry specifies the utilities (u^f_{ij}, u^w_{ij}) for f_i and w_j, respectively, when they are matched with each other.

Suppose first that there are no switching costs, meaning $\epsilon = 0$. If worker w_3 were absent, the resulting balanced market would have two stable matchings: one matching f_i with w_i and the other matching f_i with w_j, where $i = 1, 2$ and $j = 3 - i$. However, w_3 is present, introducing market imbalance, and causing the core to collapse to a unique matching.

The unique stable matching is the "diagonal" matching μ_1, where $\mu_1(i) = i$ for $i = 1, 2$. On the one hand, when w_3 is present, there are more candidate matchings for stability: $\binom{3}{2}2!$ matchings where 2 firms are paired with 2 of the 3 workers, compared to 2! in the absence of w_3. On the other hand, each candidate matching faces stricter stability restrictions due to the added competition among workers. Specifically, every firm must prefer her assigned partner to the unmatched, "outside-option" worker. For instance, matching μ_2, which pairs

[14] If single agents were prevented from forming matches, stable matchings could include unmatched firms, adding less meaningful forms of multiplicity.

f_i with w_j where $j = 3 - i$, is not stable because (f_1, w_3) forms a blocking pair: f_1 slightly prefers w_3 over w_2, that is, $1.01 = u^f_{1,3} > u^f_{1,2} = 1$. Similarly, no matching $\mu \neq \mu_1, \mu_2$ can be stable. In fact, this latter observation follows directly from the lattice structure, particularly the Rural Hospital Theorem (McVitie and Wilson, 1970) requiring the set of matched workers—here, w_1 and w_2—to remain the same across all stable matchings.

The number of ϵ-stable matchings "explodes" at the switching cost $\epsilon = 0.01$. The original stable matching, μ_1, is ϵ-stable. Additionally, μ_2 becomes ϵ-stable: although f_1 slightly prefers w_3 over w_2, the marginal benefit of switching does not outweigh the switching cost. Furthermore, matching μ_3, which pairs f_1 with w_3 and f_2 with w_1, is also ϵ-stable. Note that, in the example, there is no lattice structure on the set of ϵ-stable matchings. △

2.2 Random Markets

To assess the impact of switching costs, we consider markets with random cardinal utilities.

A *random market* is a matching market in which the utilities $\mathcal{U} = \{U^f_{i,j}, U^w_{i,j}\}_{i\in[n],j\in[n+k]}$ for matches between firms and workers are drawn independently from the exponential distribution $\text{Exp}(\lambda)$. Here, $\text{Exp}(\lambda)$ represents the exponential distribution with parameter $\lambda > 0$, which has the mean $\mathbb{E}[\text{Exp}(\lambda)] = 1/\lambda$. We use capital letters to denote random variables and small letters for their realizations.

Within this framework, we introduce relative switching costs, which represent the barriers to changing partners. The parameter λ captures the magnitude of potential match utilities. Specifically, smaller values of λ correspond to greater variability in utilities and a higher average utility of $1/\lambda$. By scaling the cost ϵ by $1/\lambda$, we define the *relative switching cost* as $\epsilon\lambda$, which measures these barriers in relation to the average match utility, independent of the absolute scale of utilities.[15]

Our benchmark model builds on the canonical random market framework with ordinal preferences. When switching costs are absent, stability hinges on ordinal comparisons alone. Thus, our setting coincides with the canonical model, where agents' strict ordinal preference lists are determined uniformly at random (Leshno, 2023)—these rankings may reflect specific geographical preferences, personal fit, family constraints, connections, and other factors.[16] Cardinal match utilities allow for more expressive preferences.

[15] Our analysis remains unchanged if agents face different switching costs and have utilities that are arbitrary positive affine transformations of the exponential distribution, provided their relative costs are consistent. Similar qualitative results hold if agents have varying relative costs.

[16] Random match utilities effectively induce strict preference rankings, as any indifferences that may arise occur with measure zero.

The exponential distribution, as highlighted in the introduction, is widely employed to model utilities across various applications. In the context of matching, it captures a natural scenario in which agents can more readily differentiate their higher-ranked partners, with such distinctions becoming less pronounced for those partners they view less favorably. In addition, this distribution brings tractability and flexibility to our model. In Section 4, we consider an alternative distribution corresponding to a complementary scenario where higher-ranked partners are less easily differentiated; this reinforces our insights and uncovers additional phenomena arising when agents are more replaceable.

In the next section, we present the main results using a sequence of random markets defined as

$$(\mathcal{F}_n, \mathcal{W}_n, \mathcal{U}_n), \quad n \geq 1,$$

where the n-th random market consists of n firms and $n+k$ workers, with $k = k(n) \geq 0$. Below, we provide some additional notation.

Asymptotic Notations We write $f(n) = O(g(n))$ as $n \to \infty$ if $\limsup_{n\to\infty} f(n)/g(n) < \infty$, and $f(n) = o(g(n))$ if $\limsup_{n\to\infty} f(n)/g(n) = 0$. In addition, $f(n) = \Omega(g(n))$ holds if $g(n) = O(f(n))$, and $f(n) = \omega(g(n))$ if $g(n) = o(f(n))$. Furthermore, $f(n) = \Theta(g(n))$ as $n \to \infty$ if both $f(n) = O(g(n))$ and $g(n) = O(f(n))$. The above limit definitions assume $g(n) > 0$ for sufficiently large n.

Finally, $f(n)$ is *polynomially large* if $f(n) = O(n^d)$ for some constant $d \geq 0$. In contrast, $f(n)$ is *super-polynomially large* if for *every* constant $d \geq 0$, $f(n) = \omega(n^d)$. Super-polynomial growth is regarded as very fast, exceeding the rate of any polynomial function.

3 Explosion of the Core

In this section, we present and discuss our main theoretical results, demonstrating how the presence of switching costs can significantly expand the set of stable outcomes in matching markets. Specifically, we identify a sharp threshold for switching costs beyond which the set of ϵ-stable matchings increases dramatically. Notably, this expansion occurs across both balanced and imbalanced markets, even when switching costs vanish as the market size grows. These results stand in contrast to the conventional belief that imbalanced, or large, markets tend to have a small number of stable matchings.

To formalize our analysis and state the results, let $S^\epsilon(n, n+k)$ denote the number of ϵ-stable matchings in a market with n firms and $n+k$ workers. For simplicity, we omit the

explicit dependence on the parameter $\lambda > 0$, whenever it does not cause confusion. Then,

Lemma 1. *Consider a random market with n firms and $(n+k)$ workers for arbitrary imbalance $k \geq 0$ and costs $\epsilon \geq 0$. Then,*

1. $\mathbb{E}\left[S^{\epsilon}(n, n+k)\right]$ *is a function of the relative switching cost* $\epsilon\lambda = \epsilon / \mathbb{E}[Exp\,(\lambda)]$*;*

2. $\mathbb{E}\left[S^{\epsilon}(n, n+k)\right]$ *is strictly increasing in* $\epsilon\lambda$*.*

Intuitively, the expected number of stable matchings is driven by the interplay between utilities and switching costs, which is effectively captured by a single parameter—relative costs. As these relative costs increase, agents become less willing to switch, leading to a greater variety of stable outcomes. We later provide an exact formula for this expectation.

In what follows, we characterize the impact of relative costs on the size of the core, i.e., the set of stable matchings. Section 3.1 sketches the arguments underlying our results.

Balanced Markets Consider first balanced markets with n firms and n workers and relative switching costs $\epsilon\lambda = \epsilon(n)\lambda(n)$, which may depend on the market size. Then,

Theorem 1. *Consider a sequence of random markets, indexed by n, with n firms and n workers, for arbitrary relative costs $\epsilon\lambda = \epsilon(n)\lambda(n) \geq 0$. Then,*

1. *if* $\epsilon\lambda = O(n^{-1}\log n)$*, then* $\mathbb{E}[S^{\epsilon}(n,n)]$ *is polynomially large;*

2. *otherwise, for* $\epsilon\lambda = \omega(n^{-1}\log n)$*,* $\mathbb{E}[S^{\epsilon}(n,n)]$ *is super-polynomially large.*

This theorem establishes a critical threshold for the relative switching cost, marking a dramatic shift in the behavior of ϵ-stable matchings. When costs decrease at least as fast as $\log n/n$ as the market size n grows, the expected number of ϵ-stable matchings can be sizable but still remains within bounds similar to those when costs are absent. Specifically, this number can exceed $e^{-1}n\log n$, the approximate expected number of stable matchings when no costs are present (Pittel, 1989), yet it remains polynomially large. However, the scenario shifts drastically when $\epsilon\lambda$ exceeds this threshold—the set of ϵ-stable matchings explodes, growing faster than any polynomial function of n.

Imbalanced Markets Importantly, our second theorem establishes that the same critical threshold applies to imbalanced markets, triggering the explosive expansion in the set of ϵ-stable matchings. Specifically, consider imbalanced markets with n firms and $n+k$ workers, where $k = k(n) \geq 1$ represents various degrees of market imbalance.

Theorem 2. *Consider a sequence of random markets, indexed by n, with n firms and $n+k$ workers, for arbitrary imbalance $k = k(n) \geq 1$ and relative costs $\epsilon\lambda = \epsilon(n)\lambda(n) \geq 0$. Then,*

1. *if $\epsilon\lambda = O(n^{-1}\log n)$, then $\mathbb{E}[S^\epsilon(n, n+k)]$ is polynomially large, uniformly for all k;*
2. *otherwise, for $\epsilon\lambda = \omega(n^{-1}\log n)$, $\mathbb{E}[S^\epsilon(n, n+k)]$ is super-polynomially large, uniformly for all k.*

Thus, the critical switching cost threshold remains consistent regardless of the degree of market imbalance.

Corollary 1. *The threshold $\epsilon\lambda = \Theta(n^{-1}\log n)$, which separates polynomial from super-polynomial growth, holds for any imbalance $k = k(n) \geq 0$.*

In many practical settings, switching costs are substantial, often being largely unrelated to market size and, in some cases, possibly increasing with it. For example, in marriage markets, the emotional toll and legal challenges of dissolving a partnership are determined by individual experiences and law, respectively, rather than market size. Contractual obligations remain binding across various labor markets (Azar and Marinescu, 2024). Additionally, psychological factors such as guilt or inertia are deeply personal and may persist, while reputational concerns and social pressures can intensify when more individuals observe and judge each other's actions (Bursztyn and Jensen, 2017).

Our results reveal that the core's explosive expansion occurs even when costs diminish with the market size.

Corollary 2. *Super-polynomial growth holds even for vanishingly small relative costs $\epsilon\lambda = \epsilon(n)\lambda(n) = o(1)$, as long as $\epsilon\lambda = \omega(n^{-1}\log n)$.*

Thus, the set of stable outcomes can be vast even in imbalanced markets with small costs. Section 4 further refines this finding by showing that the multiplicity may be substantial, involving many agents who experience significantly different outcomes across ϵ-stable matchings, even in smaller—not only large—such markets. This abundance of plausible outcomes brings indeterminacy into imbalanced markets, relative to an idealized, frictionless scenario.

These findings point to a conceptual departure from the traditional view that imbalanced, or large, markets typically have a small set of stable matchings. Absent costs, even in minimally imbalanced markets—with n firms and $n+1$ workers—the expected number of

stable matchings becomes substantially smaller compared to balanced markets;[17] and only a negligible fraction of agents have multiple stable partners (Ashlagi, Kanoria, and Leshno, 2017). As the imbalance $k = k(n)$ increases, the set of stable matchings collapses to a singleton (Pittel, 2019). This demonstrates the pivotal role that switching costs, even when small, play in dramatically expanding the set of plausible market outcomes.

Integral Formulas Central to our sharp characterization is the integral-based approach, elucidated by Knuth (1976) and developed by Pittel (1989). While previously applied in classical matching settings, we demonstrate that this approach is also powerful in the novel realm of ϵ-stable matchings. Despite the absence of structural results for these matchings, it enables us to derive tractable expressions amenable to formal analysis.

Consider a matching μ in which all firms are matched—for instance, the diagonal matching that pairs firms and workers with identical indices. Let $P^{\epsilon}(n, n+k)$ denote the probability that such a matching μ is ϵ-stable. By symmetry, this probability does not depend on the specific choice of a matching with all firms matched.

Lemma 2. *Consider a random market with n firms and $(n+k)$ workers for arbitrary imbalance $k \geq 0$ and costs $\epsilon \geq 0$. Then,*

$$P^{\epsilon}(n, n+k) = \int\limits_{\mathbf{x},\mathbf{y}\in[0,1]^n} \prod_{i\neq j}(1 - e^{-2\epsilon\lambda}x_i y_j) \prod_{l\in[n]}(1 - e^{-\epsilon\lambda}x_l)^k \, d\mathbf{x}d\mathbf{y}.$$

At an intuitive level, the integrand corresponds to the condition that no ϵ-blocking pairs exist for the considered diagonal matching.

Using the linearity of expectation and the fact that there are $\binom{n+k}{n}n!$ candidate matchings for being ϵ-stable, we get

Corollary 3. *Consider a random market with n firms and $(n+k)$ workers for arbitrary imbalance $k \geq 0$ and costs $\epsilon \geq 0$. Then,*

$$\mathbb{E}\left[S^{\epsilon}(n, n+k)\right] = \binom{n+k}{n} n! \cdot \int\limits_{\mathbf{x},\mathbf{y}\in[0,1]^n} \prod_{i\neq j}(1 - e^{-2\epsilon\lambda}x_i y_j) \prod_{l\in[n]}(1 - e^{-\epsilon\lambda}x_l)^k \, d\mathbf{x}d\mathbf{y}.$$

In particular, the expectation is a function of the relative switching cost $\epsilon\lambda = \epsilon / \mathbb{E}[Exp(\lambda)]$.

This result, generalizing Lemma 1, provides the multi-dimensional integral formula that serves as a stepping stone in our proofs; see details in the next section.

[17] Adding just one more worker already reduces the asymptotic expected number of stable matchings by a sizable factor of $\log^2 n$, bringing it down to $e^{-1}n/\log n$ (Pittel, 2019), despite the increase in candidate matchings for stability from $n!$ to $(n+1)!$.

A notable advantage of our model is its ability to yield a variety of tractable integral expressions, enabling the exploration of other important facets of matching markets, which are a focus of our current research. For instance, Lemmas 5–7 in Appendix C derive formulas that can be applied to study ϵ-stable matchings from a welfare perspective; see also Section 4 for additional discussion.[18] This tractability and flexibility establish our model as a valuable benchmark for investigating the effects of switching costs in matching markets.

3.1 Structure of Proofs

In this section, we sketch the key ideas behind the proofs, deliberately keeping the discussion less formal to avoid technical difficulties. For more formal and rigorous proofs, we refer the interested reader to the Appendix.

Transformation of Random Markets To simplify the derivation of integral formulas, we introduce an equivalent representation of random markets. In this representation, random match utilities—modeled as exponential random variables—are expressed as log-utilities over random match values. In essence, if X is uniformly distributed on $[0,1]$, then

$$\mathbb{P}\left(\log\frac{1}{X}\leq x\right)=1-e^{-x},$$

which means that $\log\frac{1}{X}$ follows the exponential distribution with rate 1. More generally, for random markets with parameter λ, we can equivalently work with $U^f_{i,j}=\frac{1}{\lambda}\log\frac{1}{X_{i,j}}$ and $U^w_{i,j}=\frac{1}{\lambda}\log\frac{1}{Y_{i,j}}$, where the random variables $\{X_{i,j},Y_{i,j}\}_{i\in[n],j\in[n+k]}$ are i.i.d. uniform on $[0,1]$. Here, $\frac{1}{X_{i,j}}$ and $\frac{1}{Y_{i,j}}$ can be interpreted as match values, over which utilities are computed.

This transformation renders the study of epsilon-stable matchings both tractable and amenable to formal analysis.

Integral Formula To provide intuition behind the integral formula in Lemma 2, consider the diagonal matching μ, with $\mu(i)=i$ for all $i\in[n]$. Additionally, using the transformation, fix the utilities of the matched agents as $U^f_{i,i}=u^f_{i,i}=\frac{1}{\lambda}\log\frac{1}{x_i}$ and $U^w_{j,j}=u^w_{j,j}=\frac{1}{\lambda}\log\frac{1}{y_j}$.

Then, firm f_i ϵ-prefers worker w_j, $j\neq i$, over her partner under μ with probability

$$\mathbb{P}\left[U^f_{i,j}>u^f_{i,i}+\epsilon\right]=\mathbb{P}\left[X_{i,j}<e^{-\epsilon\lambda}x_i\right]=e^{-\epsilon\lambda}x_i.$$

[18] A list of additional useful integral formulas is available upon request.

Similarly, non-single worker w_j ϵ-prefers firm f_i over her partner, firm f_j, with probability $e^{-\epsilon\lambda}y_j$. Importantly, the corresponding events are independent, given the fixed utilities of the matched agents.

Building on these ideas, the integral formula has an intuitive structure. The first product in the integrand corresponds to the probability that no firm f_i and non-single worker w_j form an ϵ-blocking pair. The second product represents the probability that no firm f_i and single worker w_j, where $j > n$, form an ϵ-blocking pair. Integrating over all values of x_i and y_j yields the desired formula.

This integral formula serves as the foundation of our analysis in Theorem 1 and 2.

Organization of Proofs The proofs of Theorem 1 and Theorem 2, while relying on conceptually similar techniques, differ due to the presence of an additional product in the integral formula for imbalanced markets. This product, which roughly represents competitive forces, drastically alters the dominant regions of the integration domain, necessitating careful examination and distinct treatment in each case.

Methodologically, the proof of each theorem proceeds in two steps. In the first step, we establish polynomial growth of the expected number of ϵ-stable matchings by providing an upper bound for this expectation. The second step, which requires a more intricate and refined analysis, proves super-polynomial growth by deriving a lower bound. To address the unique challenges for a sharp analysis posed by switching costs—and specifically to establish results that hold uniformly across all imbalances—we develop technical tools inspired by recent findings of Pittel (2019).

In what follows, we provide rough intuition behind Theorem 2, focusing primarily on the first step, formally proved as Proposition 3 in the Appendix. This step offers a particularly effective argument through a connection we uncover to the so-called beta functions, while still providing a hint of the more general techniques used throughout the analysis. At the end of this section, we also offer brief, informal intuition behind the substantially more involved second step, which is proved as Proposition 4.

Polynomial Growth The result for polynomial growth follows from a series of approximations and a connection to the random partition of the interval $[0, 1]$. For notational simplicity, let $p \equiv \exp(-2\epsilon\lambda)$. We start by using the first-order exponential approximation

based on the inequality $1-\xi \le \exp(-\xi)$:

$$P^{\epsilon}(n, n+k) \le \int\limits_{\mathbf{x}\in[0,1]^n} \left(\prod_{j\in[n]} \int_0^1 \exp\left(-pys_j\right)\, dy \right) \exp\left(-k\sqrt{p}s\right)\, d\mathbf{x},$$

where $s = \sum_{i\in[n]} x_i$ and $s_j = s - x_j$.

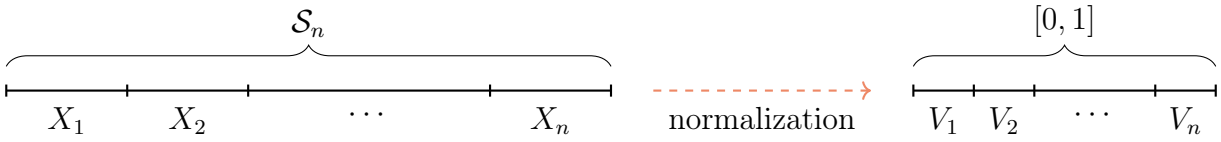
Integrating with respect to y and leveraging the idea that $s \approx s_j$, we establish that $P^{\epsilon}(n, n+k)$ is at most of order

$$\int\limits_{\mathbf{x}\in[0,1]^n} F^n(sp)e^{-k\sqrt{p}s}\, d\mathbf{x} = \int_0^n F^n(sp)e^{-k\sqrt{p}s}\, f_n(s)\, ds, \quad \text{with} \quad F(z) = \frac{1-e^{-z}}{z},$$

where the resulting integrand is a function of s only, and $f_n(s)$ is the probability density of $\mathcal{S}_n = \sum_{j\in[n]} X_j$.

In more refined approximations, additional statistics involving $\{X_j\}_{j\in[n]}$, beyond $\mathcal{S}_n$, come into play. We utilize their link to the structure of a random partition of the interval $[0,1]$—particularly to apply $f_n(s)$ in the above expression—which we briefly discuss below.

Connection to Random Partition of Unit Interval It is known that $f_n(s) \le \frac{s^{n-1}}{(n-1)!}$; see, for example, Chapter 1 in Feller (1971). Intuitively, the variables $\{X_j\}_{j\in[n]}$ normalized by $\mathcal{S}_n$ correspond to the partition $\{V_j\}_{j\in[n]}$ of the unit interval $[0,1]$, where $V_j = X_j/\mathcal{S}_n$:



It can be verified that the Jacobian of $(x_1, \ldots, x_n)$ with respect to $(s, v_1, \ldots, v_{n-1})$ is s^{n-1}, implying that the density of $(\mathcal{S}_n, V_1, \ldots, V_{n-1})$ is

$$f_n(s, v_1, \ldots, v_{n-1}) \le s^{n-1}\, \mathbb{1}\left(\sum_{j<n} v_j \le 1 \right),$$

where $\mathbb{1}(\cdot)$ denotes an indicator function. The desired inequality, $f_n(s) \le \frac{s^{n-1}}{(n-1)!}$, then follows from the observation that the lengths of the first $(n-1)$ consecutive intervals in the above partition have the joint density $(n-1)! \cdot \mathbb{1}\left(\sum_{j<n} v_j \le 1\right)$.

In our proofs, we employ various properties arising from the link to the random partition of $[0, 1]$, including the exact expression for $f_n(s)$; for further details, see Appendix A.

Polynomial Growth (Continued) Relying on this connection and the previous approximations, we derive a tractable and effective upper bound for the expectation $\mathbb{E}\left[S^{\epsilon}(n, n+k)\right]$ and establish its polynomial growth by using the properties of beta functions.

To gain intuition behind the appearance and role of the beta function in the argument, note that there are $\binom{n+k}{n} n!$ matchings that can potentially be ϵ-stable. The binomial coefficient $\binom{n+k}{n}$ can be expressed as $\frac{n+k}{nk} \frac{1}{B(n,k)}$, where

$$B(n, k) = \int_0^1 z^{n-1}(1-z)^{k-1}\, dz$$

is the beta function. Furthermore, after additional approximations and several technical steps, omitted for brevity, we find that $P^{\epsilon}(n, n+k)$ is at most of order $\frac{p^{-n}}{(n-1)!} B(n, k/\sqrt{p})$.

By collecting all the pieces, we conclude the proof. Specifically, the expected number $\mathbb{E}\left[S^{\epsilon}(n, n+k)\right]$ is at most of order $\frac{(n+k)}{k} \cdot \frac{B(n,k/\sqrt{p})}{B(n,k)} \cdot p^{-n}$. The first fraction is $O(n)$, uniformly for $k > 0$. Furthermore, the second fraction is at most 1 since $B(n, k/\sqrt{p}) \leq B(n, k)$ by the definition of the beta function. In fact, this fraction—which may be interpreted as reflecting competitive pressure—can be substantially less than 1; this could further reduce the upper bound, though such a reduction is not critical in this case. Indeed, as $\epsilon\lambda = O(n^{-1} \log n)$, the factor $p^{-n} = \exp(2\epsilon\lambda n)$ is only polynomially large, delivering the result.

Comment on Super-Polynomial Growth The result for the uniform explosion in the number of ϵ-stable matchings is significantly more involved and technical. Even setting aside approximation intricacies and a real lower bound, it is unclear without further analysis whether even the above, quite rough, upper bound—when treated as a hypothetical lower bound—would suffice for the uniform result. While p^{-n} suggests an increase in stable matchings, competitive pressure represented by $\frac{B(n,k/\sqrt{p})}{B(n,k)}$ and dependent on k could theoretically collapse the core for large imbalances k. Furthermore, if the lower bound were smaller than the upper bound by a function of k that grows, however slowly, to infinity, the uniform result would not be feasible either. To address this, our proof employs substantially more refined approximations and deeper properties stemming from the connection to the random partition of the unit interval; see Appendix B for details.

4 Computational Experiments: Integer Programming

In this section, we present simulation results for random markets that complement and refine our theoretical findings. Towards this end, we establish a link between ϵ-stable matchings and integer solutions of a specific system of linear inequalities. Leveraging this link, we illustrate that the set of ϵ-stable matchings already expands dramatically in small markets with small costs. This multiplicity translates into significantly different outcomes for many agents. We show that our benchmark model provides a conservative estimate of multiplicity compared to matching scenarios, where agents find it less easy to differentiate between their higher-ranked partners. We also discuss additional potential effects of switching costs.

Integer Programming Formulation In what follows, we show that ϵ-stable matchings can be described as integer solutions of a linear programming problem. Formally, each matching μ can be represented as an incidence vector $x \in \{0,1\}^{n\times(n+k)}$ such that

$$x_{i,j} = \begin{cases} 1 & \text{if } \mu(i) = j, \text{ i.e., } f_i \text{ is matched to } w_j, \\ 0 & \text{otherwise.} \end{cases}$$

A vector $x \in \mathbb{R}^{n\times(n+k)}$ is a matching if and only if it is an integer solution of the following system of linear inequalities:[19]

$$\begin{aligned} \sum_j x_{i,j} &\leq 1 \quad \text{for each } f_i \in \mathcal{F}, \\ \sum_i x_{i,j} &\leq 1 \quad \text{for each } w_j \in \mathcal{W}, \\ x_{i,j} &\geq 0 \quad \text{for each } (f_i, w_j) \in \mathcal{F} \times \mathcal{W}. \end{aligned} \tag{1}$$

The following lemma characterizes ϵ-stable matchings as integer solutions of a system of linear inequalities.

Lemma 3. *Let x be a matching. Then, x is ϵ-stable if and only if*

$$x_{i,j} + \sum_{k\neq j:\, u^f_{i,k}+\epsilon\geq u^f_{i,j}} x_{i,k} + \sum_{k\neq i:\, u^w_{k,j}+\epsilon\geq u^w_{i,j}} x_{k,j} \geq 1 \quad \textit{for each } (f_i, w_j) \in \mathcal{F}\times\mathcal{W}. \tag{2}$$

[19] A possibly non-integer solution $x \in \mathbb{R}^{n\times(n+k)}$ of (1) is often called a *fractional matching*. By Birkhoff's theorem, every fractional matching is a convex combination of matchings.

In particular, a vector $x \in \mathbb{R}^{n\times(n+k)}$ *is an* ϵ*-stable matching if and only if it is an integer solution of* (1) *and* (2).

This lemma and the ϵ*-stable polytope*, defined by (1) and (2), generalize those for stable matchings (Vande Vate, 1989). Intuitively, constraints (2) impose the ϵ-stability condition: for each pair (f_i, w_j), either they are matched to each other; or firm f_i is matched to some worker she prefers enough not to switch to w_j, given the cost ϵ; or w_j is matched to some firm he prefers enough not to leave for f_i. The resulting polytope is guaranteed to be non-empty, as stable matchings, which are also ϵ-stable, always exist.[20]

We use this formulation for computational experiments to examine ϵ-stable matchings.

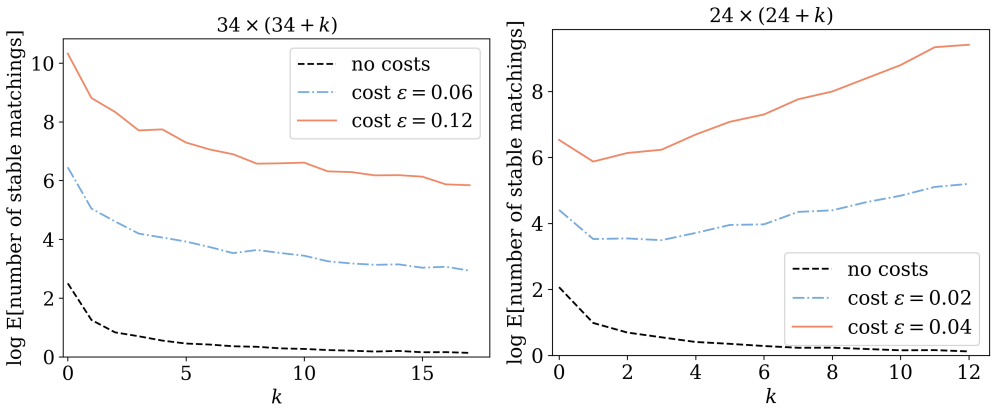
Number of Stable Matchings Figure 2 demonstrates that the set of stable matchings explodes even in small markets with small costs. We focus on markets with utilities drawn from the standard exponential distribution $\mathrm{Exp}\,(1)$, with the mean of 1—thereby making relative costs directly represented by ϵ. We also consider a distribution that corresponds to scenarios where agents find it less easy to differentiate between their higher-ranked partners, providing a complementary perspective. Specifically, we examine markets with uniformly random utilities $U[0, 2]$, maintaining the same mean of 1. Such markets result in utilities that are more evenly spread, especially for top preferences, and more compressed.

For each market size, $n\times(n+k)$, we run 1,000 simulations. While the integer programming approach enables detailed analysis, it still imposes computational restrictions. Accordingly, for the exponential case, we focus on markets with $n = 34$ firms and varying imbalances $k \geq 0$. We explore two levels of costs, $\epsilon = 0.06$ and $\epsilon = 0.12$, selected to lie below and slightly above the reference threshold, $\log n/n \approx 0.10$.[21] For the uniform distribution, constrained by computational feasibility, we examine smaller markets with $n = 24$ firms and reduced costs of $\epsilon = 0.02$ and $\epsilon = 0.04$, which already provide meaningful insights.

Notably, the multiplicity of stable matchings becomes substantial even in these very small markets with the small costs considered. In the exponential case, the set of ϵ-stable matchings is large, in contrast to the very limited number observed in a no-cost scenario. For the uniform case, greater market imbalance can even lead to an expanding set of stable outcomes—we discuss the nature of this phenomenon and the sources of multiplicity below.

[20] We show in Example 2 in Appendix D that extreme points of the ϵ-stable polytope may be non-integer. This contrasts with the standard stable polytope, which is always the convex hull of stable matchings (Vande Vate, 1989; Roth, Rothblum, and Vande Vate, 1993; Teo and Sethuraman, 1998).

[21] Similar qualitative results hold for smaller costs.



(a) Exponential (b) Uniform

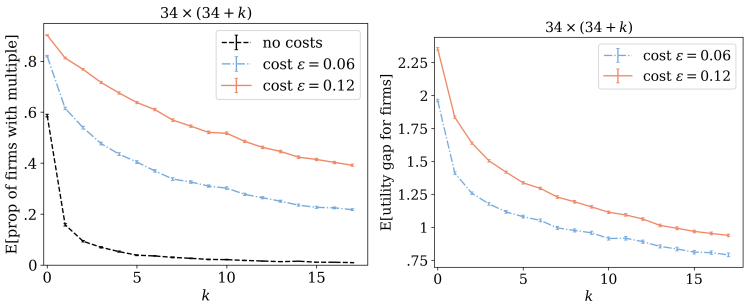
Figure 2: log [Number of ϵ-stable matchings]

Multiplicity Patterns We demonstrate that the substantial multiplicity manifests in many agents experiencing significantly different outcomes across ϵ-stable matchings.

Figure 3 presents the multiplicity features for our benchmark exponential model. A significant fraction of firms have multiple ϵ-stable partners across different imbalances, even sizable ones; see Panel (3a). This contrasts with a frictionless scenario, where even minimal imbalances yield an essentially unique stable matching with a negligible proportion of agents having multiple partners (Ashlagi, Kanoria, and Leshno, 2017). Appendix E illustrates similar patterns for workers—it further shows that ϵ-stable matchings may involve many mismatched participants compared to the essentially unique stable matching.

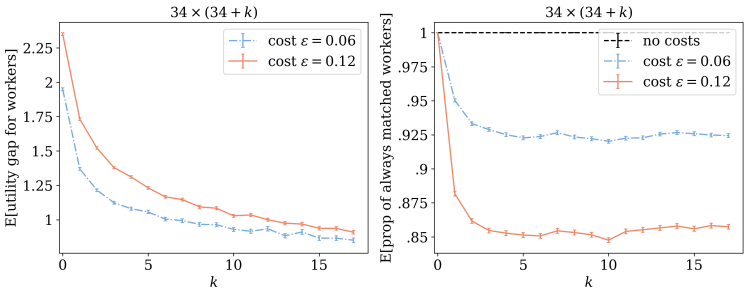
Agents not only have multiple ϵ-stable partners, but these partners can differ significantly. Panel (3b) shows that the average utility gap for firms with multiple ϵ-stable partners is substantial, especially compared to the given small costs; the no-cost case is omitted, as only a negligible fraction of agents have multiple partners.[22] For workers, being (un)matched in one ϵ-stable matching does not guarantee being (un)matched in another, unlike in the traditional setting where the Rural Hospital Theorem (McVitie and Wilson, 1970) ensures a consistent set of matched agents across all stable matchings. Panel (3c) highlights significant utility gaps for consistently matched workers that have multiple ϵ-stable partners. Panel (3d) shows that, while many agents remain consistently matched, the fraction of such consistently matched workers—relative to n—is still below 1. Relatedly, Appendix E reveals that workers

[22]The utility gap is defined as the difference between the highest and lowest utility an agent obtains from their ϵ-stable partners.



(a) Proportion of firms with multiple partners

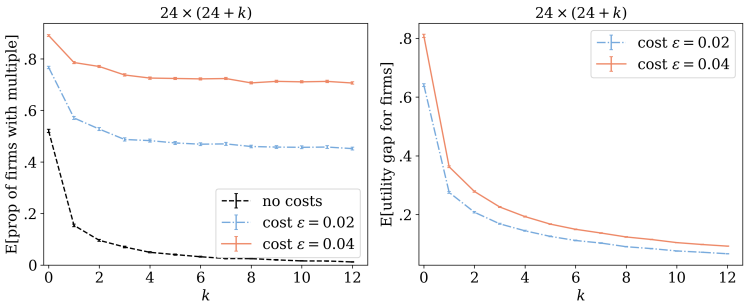
(b) Utility gap* for firms



(c) Utility gap* for always matched workers

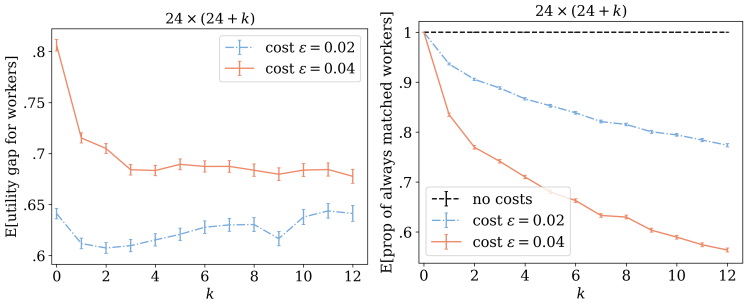
(d) Proportion of always matched workers

Figure 3: Random $34 \times (34 + k)$ markets with exponential utilities



(a) Proportion of firms with multiple partners

(b) Utility gap* for firms



(c) Utility gap* for always matched workers

(d) Proportion of always matched workers

Figure 4: Random $24 \times (24 + k)$ markets with uniform utilities

who would be unmatched in the standard setting can often find ϵ-stable partners.[23]

In a certain sense, our benchmark model provides a conservative estimate of multiplicity. In the uniform case, firms find it harder to differentiate between their higher-ranked workers, making these workers substantially more replaceable. This limited differentiation leads to a large fraction of inconsistently (un)matched workers while producing smaller utility gaps for firms with multiple ϵ-stable partners; see Panels (4d) and (4b) of Figure 4. Worker replaceability becomes more pronounced with greater imbalances, driving the high proportion of agents with multiple ϵ-stable partners, as shown in Panel (4a). It also contributes to the growing number of ϵ-stable matchings in Figure 2. Even those "lucky" workers matched in all ϵ-stable matchings—and having multiple ϵ-stable partners—experience large utility gaps, as illustrated in Panel (4c).

We provide additional computational results in Appendix E.

Additional Impacts of Switching Costs We conclude by mentioning one current direction in our research. Our model allows us to derive integral formulas that can be used to study ϵ-stable matchings from additional perspectives, including welfare; see Appendix C.

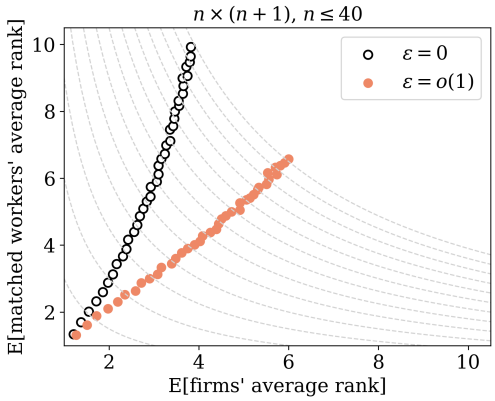


Figure 5: Firms' and matched workers' average ranks in the ϵ-stable matching that minimizes the matched workers' average rank

Figure 5 suggests that, with switching costs, the average welfare of firms and matched workers may be nuanced. Without switching costs, even minimal imbalances not only cause the core to collapse but also significantly benefit the short side of the market (Ashlagi,

[23]For traditional stability, the theoretical study of certain aspects, such as the multiplicity of stable partners, generally relies on its well-defined lattice structure and related properties, including rotations (Irving and Leather, 1986) and rejection chains in McVitie and Wilson (1971)'s algorithm. No analogous structure has yet been identified for ϵ-stable matchings.

Kanoria, and Leshno, 2017). Specifically, even in markets with n firms and $n+1$ workers, the firms' average rank of their paired workers is $\log n(1+o(1))$, while the matched workers' average rank of their paired firms is $\frac{n}{\log n}(1+o(1))$.[24] Suppose the costs are given by $\epsilon(n) = n^{-1}\log n \cdot w(n)$, where $w(n)$ is of order $n^{1/2}$, so that $\epsilon(n) = o(1)$.[25] Figure 5 indicates that even such small costs may potentially lead to significant welfare gains for matched workers.[26]

5 Conclusions

This paper studies the impacts of switching costs on outcomes in matching markets. We introduce a new notion of stability specifically tailored to account for switching costs and examine its implications. Switching costs lead to an abundance of stable matchings, even in large imbalanced markets, and even when these costs are very small. We provide a sharp characterization of the threshold switching cost that triggers an explosion in the number of stable matchings. These findings suggest that switching costs may serve as a valuable tool in market design, offering significant leverage for policy interventions. Our framework can be used to study further effects of switching costs, across a variety of matching environments that have been traditionally investigated only through the lens of classical stability.

[24] For firm f_i, the rank of worker w_j in f_i's preferences is the number of workers that are at least as good as w_j. Smaller ranks correspond to better partners, and f_i's most preferred worker has rank 1. For workers, the rank of firm f_i in w_j's preferences is defined symmetrically.

[25] Formally, $w(n) = n^{1/2}/6$, where the normalization ensures that costs are small within the considered region of market sizes. Similar patterns emerge for other $w(n)$. As in the previous exercises, we simulate 1,000 markets per market size, with utilities drawn from the standard exponential distribution.

[26] The figure further suggests that, for ϵ-stable matchings, the product of the average ranks for firms and matched workers may be close to n. This resembles the "law of hyperbola" established for traditional stable matchings in balanced markets (Pittel, 1992).

Appendix A. Proof of Lemma 2 and Theorem 1

In this section, we prove Lemma 2 and Theorem 1 stated in Section 3.

Transformation of Random Markets A random market with i.i.d. Exp (λ) match utilities can be equivalently represented as one where all agents have log-utilities over match values, with these match values drawn i.i.d. from an inverse uniform distribution. Formally, we can consider $U^f_{i,j} = \frac{1}{\lambda}\log\frac{1}{X_{i,j}}$ and $U^w_{i,j} = \frac{1}{\lambda}\log\frac{1}{Y_{i,j}}$, where $\{X_{i,j}, Y_{i,j}\}_{i\in[n],j\in[n+k]}$ are $2n(n+k)$ i.i.d. random variables uniformly distributed on $[0,1]$. This transformation is valid because if U $[0,1]$ is uniform on $[0,1]$, then $\frac{1}{\lambda}\log\frac{1}{\mathrm{U}[0,1]} \sim \mathrm{Exp}\,(\lambda)$. We employ this equivalent representation to prove Lemma 2 and derive additional integral formulas in Appendix C, with their asymptotic analysis being an active part of our ongoing research.

Proof of Lemma 2. Consider first the balanced market with $k = 0$. By symmetry, it suffices to evaluate $P^\epsilon(n,n)$, the probability that the diagonal matching μ is ϵ-stable. Conditioned on $X_{i,i} = x_i$ and $Y_{j,j} = y_j$, the $n(n-1)$ events

$$\begin{aligned} B_{i,j} &= \left\{U^f_{i,j} > U^f_{i,i} + \epsilon,\ U^w_{i,j} > U^w_{j,j} + \epsilon\right\} \\ &= \left\{\frac{1}{\lambda}\log\frac{1}{X_{i,j}} > \frac{1}{\lambda}\log\frac{1}{X_{i,i}} + \epsilon,\ \frac{1}{\lambda}\log\frac{1}{Y_{i,j}} > \frac{1}{\lambda}\log\frac{1}{Y_{j,j}} + \epsilon\right\}, \quad i \neq j, \end{aligned}$$

are independent, and denoting the conditioning in question by $|\circ)$, we have

$$\mathbb{P}(B_{i,j}|\circ) = \mathbb{P}\left(X_{i,j} \le e^{-\epsilon\lambda}x_i,\ Y_{i,j} \le e^{-\epsilon\lambda}y_j\right) = e^{-2\epsilon\lambda}x_i y_j.$$

Therefore,

$$P^\epsilon(n,n) = \int\limits_{\mathbf{x},\mathbf{y}\in[0,1]^n} \prod_{i\neq j}\left(1 - \mathbb{P}(B_{i,j}|\circ)\right)\, d\mathbf{x}d\mathbf{y} = \int\limits_{\mathbf{x},\mathbf{y}\in[0,1]^n} \prod_{i\neq j}(1 - e^{-2\epsilon\lambda}x_i y_j)\, d\mathbf{x}d\mathbf{y}.$$

For $k > 0$, introduce events $B_{i,j}$, $i \in [n]$, $j \in [n+1, n+k]$, defined as follows:

$$B_{i,j} = \left\{U^f_{i,j} > U^f_{i,i} + \epsilon\right\} = \left\{\frac{1}{\lambda}\log\frac{1}{X_{i,j}} > \frac{1}{\lambda}\log\frac{1}{X_{i,i}} + \epsilon\right\}.$$

These events are conditionally independent on each other and on $B_{i,j}$, $i,j \in [n]$, and

$$\mathbb{P}(B_{i,j}|\circ) = e^{-\epsilon\lambda}x_i, \quad j \in [n+1, n+k].$$

Therefore, the probability that μ matching diagonally the first n firms and the first n workers is ϵ-stable is given by

$$P^{\epsilon}(n, n+k) = \int_{\mathbf{x},\mathbf{y}\in[0,1]^n} \prod_{i\neq j}(1 - e^{-2\epsilon\lambda}x_i y_j) \prod_{l\in[n]}(1 - e^{-\epsilon\lambda}x_l)^k \, d\mathbf{x}d\mathbf{y}. \quad \blacksquare$$

Proof of Corollary 3. There are $\binom{n+k}{n}n!$ matchings that can potentially be ϵ-stable. Each such matching is ϵ-stable with the same probability, $P^{\epsilon}(n, n+k)$. By the linearity of expectation, $\mathbb{E}\left[S^{\epsilon}(n, n+k)\right] = \binom{n+k}{n}n! \cdot P^{\epsilon}(n, n+k)$, as desired. $\blacksquare$

In the remainder of this appendix, we use the integral formula derived in Corollary 3 to prove Theorem 1 in two steps. Specifically, Proposition 1 establishes the first part of the theorem by deriving an upper bound for $\mathbb{E}\left[S^{\epsilon}(n, n)\right]$. Proposition 2, in turn, proves the second part of the theorem by deriving a lower bound for this expected number. In the course of our asymptotic analysis in the proof of Theorem 1, as well as in the proof of Theorem 2 from Appendix B, we use several results on the properties of a random partition of the interval $[0, 1]$, which are stated below.

Preliminaries Let $X_1, \ldots, X_n$ be independent random variables each distributed uniformly on $[0, 1]$. Denote

$$\mathcal{S}_n = \sum_{j\in[n]} X_j \quad \text{and} \quad V_j = \frac{X_j}{\mathcal{S}_n}, \quad j \in [n],$$

so that $V_j \in [0, 1]$ and $\sum_{j\in[n]} V_j = 1$. Let $f_n(s, v_1, \ldots, v_{n-1})$ denote the density of $(\mathcal{S}_n, V_1, \ldots V_{n-1})$. It is straightforward to verify that the Jacobian of $(x_1, \ldots, x_n)$ with respect to $(s, v_1, \ldots, v_{n-1})$ is s^{n-1}, implying that over the relevant support region

$$f_n(s, v_1, \ldots, v_{n-1}) = s^{n-1} \mathbb{1}\left(\sum_{j<n} v_j \leq 1 \text{ and } \max_{j\leq n} v_j \leq \frac{1}{s}\right), \tag{3}$$

where $v_n = 1 - \sum_{j<n} v_j$ and $\mathbb{1}(\cdot)$ is an indicator function.

Furthermore, it is known that the joint density $g_n(v_1, \ldots, v_{n-1})$ of $(L_1, \ldots, L_{n-1})$—the lengths of the first $(n-1)$ consecutive intervals obtained by selecting $(n-1)$ points independently and uniformly at random in $[0, 1]$—follows a Dirichlet distribution with parameters

$\underbrace{(1,1,\ldots,1)}_{n-1 \text{ times}}$ and has the form $(n-1)! \cdot \mathbb{1}\left(\sum_{j<n} v_j \leq 1\right)$; for instance, see Chapter 1 in Feller (1971) for additional details.

Therefore, by integrating (3) over $(v_1, \ldots, v_{n-1})$, we obtain the density $f_n(s)$ of $\mathcal{S}_n$:

$$f_n(s) = \frac{s^{n-1}}{(n-1)!} \cdot \mathbb{P}\left(\max_{j\in[n]} L_j \leq \frac{1}{s}\right). \tag{4}$$

In addition, by dropping the constraint $\max_{j\leq n} v_j \leq \frac{1}{s}$ in (3), we get

$$f_n(s, v_1, \ldots, v_{n-1}) \leq \frac{s^{n-1}}{(n-1)!} \cdot g_n(v_1, \ldots, v_{n-1}), \tag{5}$$

where $g_n(v_1, \ldots, v_{n-1})$ is the previously defined density of $(L_1, \ldots, L_{n-1})$.

In our proofs, we use the connection to the random partition of $[0,1]$, as well as the asymptotic properties of

$$L_n^+ \equiv \max_{j\in[n]} L_j \quad \text{and} \quad T_n \equiv \sum_{j\in[n]} L_j^2,$$

which are stated in the lemma below.

Lemma 4. *In probability,*

$$\lim_{n\to\infty} \frac{L_n^+}{n^{-1}\log n} = 1 \quad \textit{and} \quad \lim_{n\to\infty} n\, T_n = 2.$$

Remark. The proof of this lemma is based on the classic result that $\{L_j : j \in [n]\}$ has the same distribution as $\left\{W_j / \sum_{i\in[n]} W_i : j \in [n]\right\}$, where $W_1, \ldots, W_n$ are independent exponentials with parameter 1; see Pittel (1989) for further references and the proof of Lemma 4.

Proposition 1. *If $\epsilon\lambda = O(n^{-1}\log n)$, then $\mathbb{E}[S^\epsilon(n,n)]$ is polynomially large.*

Proof. As shown before, $\mathbb{E}[S^\epsilon(n,n)] = n! P^\epsilon(n,n)$, where

$$P^\epsilon(n,n) = \int_{\mathbf{x},\mathbf{y}\in[0,1]^n} \prod_{i\neq j} (1 - p x_i y_j)\, d\mathbf{x} d\mathbf{y} \quad \text{and} \quad p \equiv \exp(-2\epsilon\lambda).$$

We prove the desired result by analyzing an upper bound for $P^\epsilon(n,n)$.

Introduce $s \equiv \sum_{i\in[n]} x_i$ and $s_j \equiv \sum_{i\neq j} x_i$. Since $1-\xi \le \exp(-\xi)$, we have

$$P^{\epsilon}(n,n) \le \int\limits_{\mathbf{x}\in[0,1]^n} \left(\prod_{j\in[n]} \int_0^1 \exp\left(-pys_j\right)\, dy \right) d\mathbf{x}.$$

Integrating with respect to $y \in [0,1]$, we obtain

$$P^{\epsilon}(n,n) \le \int\limits_{\mathbf{x}\in[0,1]^n} \prod_{j\in[n]} F(s_j p)\, d\mathbf{x}, \quad \text{where} \quad F(z) \equiv \frac{1-e^{-z}}{z}.$$

To simplify the above estimate, note that for $z > 0$,

$$0 > (\log F(z))' = \frac{1}{e^z - 1} - \frac{1}{z},$$

which is asymptotic to $-\frac{1}{z}$ as $z \to \infty$, and converges to $-\frac{1}{2}$ as $z \downarrow 0$. In fact, it can be readily shown that $0 < \sup_{z>0}(z+1)\left|(\log F(z))'\right| \le 2$, so that we have

$$\left|(\log F(z))'\right| \le \frac{2}{z+1}.$$

Therefore, since $s_j = s - x_j$,

$$\begin{aligned}
\log F(s_j p) &= \log F(sp) - \int\limits_{s-x_j}^{s} (\log F(zp))'_z\, dz \le \log F(sp) + \int\limits_{s-x_j}^{s} p\, \frac{2}{pz+1}\, dz \\
&\le \log F(sp) + \int\limits_{s-x_j}^{s} p\, \frac{2}{p(s-x_j)+1}\, dz \\
&\le \log F(sp) + \frac{2px_j}{(s-x_j)p+1} \le \log F(sp) + \frac{2x_j}{s}.
\end{aligned}$$

Thus, by the definition of s,

$$\prod_{j\in[n]} F(s_j p) \le e^2 F^n(sp).$$

Next, by changing variables and applying equation (4),

$$P^{\epsilon}(n,n) \le \int_{\mathbf{x}\in[0,1]^n} e^2 F^n(sp)\, d\mathbf{x} \le e^2 \int_0^n F^n(sp)\, \frac{s^{n-1}}{(n-1)!}\, ds$$
$$= e^2 \int_0^n \left(\frac{1-e^{-sp}}{sp}\right)^n \frac{s^{n-1}}{(n-1)!}\, ds = \int_0^n \frac{e^2\, p^{-n}}{(n-1)!} \frac{(1-e^{-sp})^n}{s}\, ds.$$

We fix $s(n) \equiv \frac{\log n}{2p}$ and break the last integral into two parts, $\int_1$ for $s \le s(n)$ and $\int_2$ for $s > s(n)$. In what follows, we provide upper bounds for these two integrals.

Integral 1. Note that

$$\int_1 = \frac{e^2\, p^{-n}}{(n-1)!} \int_0^{ps(n)} \exp(\Phi_n(\eta))\, d\eta,$$

where $\Phi_n(\eta) \equiv n \log(1-e^{-\eta}) - \log \eta$.

It is straightforward to verify that $\Phi_n(\eta)$ is increasing and concave on $[0, ps(n)]$ for sufficiently large n. Indeed, since $\frac{e^\eta - 1}{\eta}$ is increasing for $\eta > 0$,

$$\Phi_n'(\eta) = \frac{n}{e^\eta - 1} - \frac{1}{\eta} = \frac{n}{e^\eta - 1}\left(1 - \frac{e^\eta - 1}{n\eta}\right)$$
$$\ge \frac{n}{e^\eta - 1}\left(1 - \frac{e^{ps(n)} - 1}{nps(n)}\right) = \frac{n}{e^\eta - 1}\left(1 - O(n^{-1/2})\right) > 0$$

if n is sufficiently large, where the last equality follows from $ps(n) = \frac{\log n}{2}$. Similarly, since $\frac{e^{\eta/2} - e^{-\eta/2}}{\eta}$ is increasing for $\eta > 0$,

$$\Phi_n''(\eta) = -\frac{n}{\left(e^{\eta/2} - e^{-\eta/2}\right)^2} + \frac{1}{\eta^2} = -\frac{n}{\left(e^{\eta/2} - e^{-\eta/2}\right)^2}\left(1 - \frac{(e^{\eta/2} - e^{-\eta/2})^2}{n\eta^2}\right)$$
$$\le -\frac{n}{\left(e^{\eta/2} - e^{-\eta/2}\right)^2}\left(1 - \frac{\left(e^{ps(n)} - e^{-ps(n)}\right)^2}{n(ps(n))^2}\right)$$
$$= -\frac{n}{(e^{\eta/2} - e^{-\eta/2})^2}\left(1 - O(\log^{-2} n)\right) < 0$$

for sufficiently large n.

As a result, by using the tangent line inequality $\Phi_n(\eta) \le \Phi_n(ps(n)) + \Phi_n'(ps(n))(\eta - ps(n))$ for $\eta \in [0, ps(n)]$, we obtain

$$\int_1 = \frac{e^2\, p^{-n}}{(n-1)!} \int_0^{ps(n)} e^{\Phi_n(\eta)}\, d\eta \le \frac{e^2\, p^{-n}}{(n-1)!} e^{\Phi_n(ps(n))} \int_0^{ps(n)} \exp\left(\Phi_n'(ps(n))(\eta - ps(n))\right)\, d\eta$$
$$\le \frac{e^2\, p^{-n}}{(n-1)!} \frac{e^{\Phi_n(ps(n))}}{\Phi_n'(ps(n))}.$$

To sum up, the contribution of $n! \int_1$ to $\mathbb{E}[S^\epsilon(n,n)]$ is

$$n! \int_1 \le n! \frac{e^2\, p^{-n}}{(n-1)!} \frac{(1-e^{-ps(n)})^n}{ps(n)} \frac{e^{ps(n)}-1}{n} \left(1 - \frac{e^{ps(n)}-1}{nps(n)}\right)^{-1}$$
$$= O(p^{-n} n^{1/2}) = \exp\left(O(\log n)\right)$$

since $p = e^{-2\epsilon\lambda}$ and $\epsilon\lambda = O(n^{-1}\log n)$.

Integral 2. We now turn attention to the second integral

$$\int_2 = \int_{s(n)}^n \frac{e^2\, p^{-n}}{(n-1)!} \frac{(1-e^{-sp})^n}{s}\, ds,$$

which can be bounded as follows:

$$\int_2 \le \frac{e^2\, p^{-n}}{(n-1)!} \int_{s(n)}^n \frac{1}{s}\, ds \le \frac{e^2\, p^{-n}}{(n-1)!} \log n.$$

Consequently, the contribution of $n! \int_2$ to $\mathbb{E}[S^\epsilon(n,n)]$ is

$$n! \int_2 = O\left(np^{-n}\log n\right) = \exp(O(\log n))$$

because $p = e^{-2\epsilon\lambda}$ and $\epsilon\lambda = O(n^{-1}\log n)$. ∎

Proposition 2. *If $\epsilon\lambda = \omega(n^{-1}\log n)$, i.e., $\epsilon\lambda \gg n^{-1}\log n$, then $\mathbb{E}[S^\epsilon(n,n)]$ is super-polynomially large.*

Proof. Since $\mathbb{E}[S^\epsilon(n,n)]$ is increasing in $\epsilon\lambda$, and $\epsilon\lambda = \omega(n^{-1}\log n)$, it is sufficient to focus only on $p = e^{-2\epsilon\lambda} = 1 - o(1)$ in our calculations below.

Define the following variables:

$$s \equiv \sum_{i\in[n]} x_i, \quad s_j \equiv \sum_{i\neq j} x_i = s - x_j, \quad t \equiv \sum_{i\in[n]} x_i^2, \quad \text{and} \quad t_j \equiv \sum_{i\neq j} x_i^2 = t - x_j^2.$$

Restrict attention to the set D of all non-negative $x = (x_1, x_2, \ldots, x_n)$ such that

$$\begin{aligned} &\frac{n}{\log n \cdot \log(ns_1(n))} \le s \le \frac{n}{\log n \cdot \log(ns_2(n))}, \\ &s^{-1}x_i \le \frac{\log(ns_2(n))}{n}, \quad i \in [n], \\ &s^{-2}t \le \frac{3}{n}, \end{aligned} \tag{6}$$

where $s_1(n) > s_2(n)$, with $s_2(n) \to \infty$ as $n \to \infty$, are to be defined later. (The two bottom inequalities are essentially dictated by the connection between the fractions $s^{-1}x_i$ and the lengths L_i of random sub-intervals of $[0,1]$, as well as by Lemma 4.) According to the first two conditions, we have

$$x_i \le s\frac{\log(ns_2(n))}{n} = (\log n)^{-1} < 1$$

for sufficiently large n, so that $D \subset [0,1]^n$. Consequently, for such large n,

$$P^{\epsilon}(n) \ge \tilde{P}^{\epsilon}(n) \equiv \int_{\mathbf{x}\in D} \left(\prod_{j\in[n]} \left(\int_0^1 \prod_{i\neq j} (1 - px_iy_j)\, dy_j \right) \right) d\mathbf{x}.$$

Since $1 - \alpha = \exp(-\alpha - \alpha^2(1 + O(\alpha))/2)$ as $\alpha \to 0$, for each fixed j, we have

$$\begin{aligned} \prod_{i\in[n], i\neq j} (1 - px_iy_j) &= \exp\left(-py_js_j - p^2y_j^2t_j\frac{1 + O\left((\log n)^{-1}\right)}{2} \right) \\ &\ge \exp\left(-py_js_j - \frac{2p^2y_j^2t_j}{3} \right) \\ &\ge \exp\left(-py_js - \frac{2p^2y_j^2t}{3} \right). \end{aligned}$$

Therefore,

$$\begin{aligned} \prod_{j\in[n]} \left(\int_0^1 \prod_{i\neq j} (1 - px_iy_j)\, dy_j \right) &\ge \left(\int_0^1 \exp\left(-pys - \frac{2p^2y^2t}{3} \right) dy \right)^n \\ &= (ps)^{-n} \left(\int_0^{ps} \exp\left(-\eta - \frac{2\eta^2 t}{3s^2} \right) d\eta \right)^n \\ &\ge (ps)^{-n} \left(\int_0^{ps} \exp\left(-\eta - \frac{2\eta^2}{n} \right) d\eta \right)^n \end{aligned}$$

since $t/s^2 \le \frac{3}{n}$ for $\mathbf{x} \in D$.

By noting that $\int_0^{ps} e^{-\eta}\, d\eta = 1 - e^{-ps}$ and employing Jensen's inequality for the exponential function, we can bound

$$\begin{aligned}\int_0^{ps} \exp\left(-\eta - \frac{2\eta^2}{n}\right) d\eta &= (1-e^{-ps}) \int_0^{ps} \exp\left(-\frac{2\eta^2}{n}\right) \cdot \frac{e^{-\eta}}{1-e^{-ps}}\, d\eta \\ &\ge (1-e^{-ps}) \exp\left(-\frac{2/n}{1-e^{-ps}} \int_0^{ps} e^{-\eta}\eta^2\, d\eta\right) \\ &= (1-e^{-ps}) \exp\left(-\frac{4}{n} \cdot \frac{e^{ps} - (1+ps+(ps)^2/2)}{e^{ps}-1}\right).\end{aligned}$$

Since

$$\frac{e^{ps} - (1+ps+(ps)^2/2)}{e^{ps}-1} \le 1,$$

we obtain

$$\int_0^{ps} \exp\left(-\eta - \frac{2\eta^2}{n}\right) d\eta \ge (1-e^{-ps})e^{-4/n}.$$

Consequently,

$$\prod_{j\in[n]} \left(\int_0^1 \prod_{i\ne j} (1 - p x_i y_j)\, dy_j \right) \ge e^{-4} \left(\frac{1-e^{-ps}}{ps}\right)^n,$$

implying that

$$\tilde{P}^{\epsilon}(n) \ge e^{-4}\, p^{-n} \int_{\mathbf{x}\in D} \left(\frac{1-e^{-ps}}{s}\right)^n d\mathbf{x}.$$

We then switch to new variables $(u, v_1, v_2, \ldots, v_{n-1})$ defined as follows:

$$u \equiv \sum_{i\in[n]} x_i = s \qquad \text{and} \qquad v_i \equiv s^{-1} x_i, \quad i \in [n-1].$$

We also define $v_n \equiv s^{-1}x_n = 1 - \sum_{i\in[n-1]} v_i$. The conditions in (6) then become

$$\begin{aligned}&\frac{n}{\log n \cdot \log(n s_1(n))} \le u \le \frac{n}{\log n \cdot \log(n s_2(n))}, \\ &v_i \le \frac{\log(n s_2(n))}{n}, \quad i \in [n], \\ &\sum_{i\in[n]} v_i^2 \le \frac{3}{n}.\end{aligned} \tag{7}$$

Importantly, for these new variables, the resulting region is the direct product of the range of u and the range of $\mathbf{v}$. By using equation (3), we then get

$$\tilde{P}^{\epsilon}(n) \geq \frac{e^{-4}\, p^{-n}}{(n-1)!} \left(\int_{\frac{n}{\log n \cdot \log(n s_1(n))}}^{\frac{n}{\log n \cdot \log(n s_2(n))}} \frac{(1-e^{-pu})^n}{u}\, du \right) \times \mathbb{P}\left(L_n^+ \leq \frac{\log(n s_2(n))}{n} \quad \text{and} \quad T_n \leq \frac{3}{n} \right),$$

where $L_n^+ \equiv \max_{i \in [n]} L_i$ and $T_n \equiv \sum_{i \in [n]} L_i^2$. By Lemma 4, the probability of the event on the right-hand side tends to one as $n \to \infty$.

Choose $s_1(n) = o(\log n)$ and $s_2(n) = \frac{s_1(n)}{2}$. Then,

$$\int_{\frac{n}{\log n \cdot \log(n s_1(n))}}^{\frac{n}{\log n \cdot \log(n s_2(n))}} \frac{1}{u}\, du = \log\left(\frac{\log(n s_1(n))}{\log(n s_2(n))} \right)$$

is asymptotic to $\frac{\log 2}{\log n}$. Furthermore, since $p = 1 - o(1)$, for the considered range of u, we have

$$pu \geq \frac{pn}{\log n \cdot \log(n s_1(n))} \geq \frac{n}{2(\log n)^2},$$

so that

$$(1-e^{-pu})^n = \left(1 - \exp\left(-\frac{n}{2(\log n)^2} \right) \right)^n = 1 - o(1).$$

To conclude,

$$\tilde{P}^{\epsilon}(n) \geq \frac{e^{-4}\, p^{-n}}{(n-1)!} \frac{\log 2}{\log n} (1 - o(1)),$$

and thus

$$\mathbb{E}[S^{\epsilon}(n,n)] \geq n!\, \tilde{P}^{\epsilon}(n) \geq \frac{n\, e^{-4}\, p^{-n} \log 2}{\log n} (1 - o(1)) = \exp(\omega(\log n))$$

since $p = e^{-2\epsilon\lambda}$ and $\epsilon\lambda = \omega(n^{-1} \log n)$. ∎

Appendix B. Proof of Theorem 2

In this section, we prove Theorem 2 presented in Section 3. Proposition 3 establishes the first part of the theorem by deriving an upper bound for $\mathbb{E}[S^{\epsilon}(n, n+k)]$, while Proposition 4 handles the second part by deriving a lower bound for this expectation.

Proposition 3. *If $\epsilon\lambda = O(n^{-1}\log n)$, then $\mathbb{E}[S^{\epsilon}(n, n+k)]$ is polynomially large, uniformly for all $k \geq 1$.*

Proof. As demonstrated in Corollary 3, $\mathbb{E}[S^{\epsilon}(n, n+k)] = \binom{n+k}{n} n! \cdot P^{\epsilon}(n, n+k)$, where

$$P^{\epsilon}(n, n+k) = \int\limits_{\mathbf{x},\mathbf{y}\in[0,1]^n} \prod_{i\neq j}(1 - p x_i y_j) \prod_{l\in[n]} (1 - \sqrt{p}x_l)^k \, d\mathbf{x}d\mathbf{y} \quad \text{and} \quad p \equiv \exp(-2\epsilon\lambda).$$

We prove the desired result by analyzing an upper bound for $P^{\epsilon}(n, n+k)$.

Since $1 - \xi \leq \exp(-\xi)$, we get

$$P^{\epsilon}(n, n+k) \leq \int\limits_{\mathbf{x}\in[0,1]^n} \left(\prod_{j\in[n]} \int_0^1 \exp\left(-p y s_j\right) \, dy \right) \exp\left(-k\sqrt{p}s\right) \, d\mathbf{x},$$

where

$$s \equiv \sum_{i\in[n]} x_i \quad \text{and} \quad s_j \equiv \sum_{i\neq j} x_i = s - x_j.$$

By integrating with respect to $y \in [0, 1]$, we obtain

$$P^{\epsilon}(n, n+k) \leq \int\limits_{\mathbf{x}\in[0,1]^n} \left(\prod_{j\in[n]} F(s_j p) \right) \exp\left(-k\sqrt{p}s\right) \, d\mathbf{x}, \quad \text{where} \quad F(z) \equiv \frac{1 - e^{-z}}{z}.$$

As in the proof of Proposition 1,

$$\prod_{j\in[n]} F(s_j p) \leq e^2 F^n(sp).$$

Therefore, by changing variables and using equation (4), we get

$$
\begin{aligned}
P^{\epsilon}(n, n+k) &\leq \int_{\mathbf{x}\in[0,1]^n} e^2 F^n(sp) e^{-k\sqrt{p}s}\, d\mathbf{x} \leq e^2 \int_0^n F^n(sp) e^{-k\sqrt{p}s} \frac{s^{n-1}}{(n-1)!}\, ds \\
&= e^2 \int_0^\infty \left(\frac{1-e^{-sp}}{sp}\right)^n e^{-k\sqrt{p}s} \frac{s^{n-1}}{(n-1)!}\, ds = \frac{e^2\, p^{-n}}{(n-1)!} \int_0^\infty \frac{1}{s}(1-e^{-sp})^n e^{-k\sqrt{p}s}\, ds \\
&= \frac{e^2\, p^{-n}}{(n-1)!} \int_0^\infty \frac{1}{\eta}(1-e^{-\eta})^n \exp\left(-\frac{k\eta}{\sqrt{p}}\right) d\eta \\
&\leq \frac{e^2\, p^{-n}}{(n-1)!} \int_0^\infty (1-e^{-\eta})^{n-1} \exp\left(-\frac{k\eta}{\sqrt{p}}\right) d\eta \\
&= \frac{e^2\, p^{-n}}{(n-1)!} \int_0^1 (1-z)^{n-1} z^{\frac{k}{\sqrt{p}}-1}\, dz = \frac{e^2\, p^{-n}}{(n-1)!} B(n, k/\sqrt{p}),
\end{aligned}
$$

where $B(a_1, a_2) \equiv \int_0^1 z^{a_1-1}(1-z)^{a_2-1}\, dz$ is the beta function, also known as the Euler integral of the first kind.

The beta function is closely related to the gamma function:

$$
B(a_1, a_2) = \frac{\Gamma(a_1)\Gamma(a_2)}{\Gamma(a_1+a_2)}, \quad \text{where} \quad \Gamma(a) \equiv \int_0^\infty z^{a-1} e^{-z}\, dz,
$$

which is known as the Euler integral of the second kind. Furthermore, the beta function is also connected to the binomial coefficients

$$
\binom{n+k}{n} = \frac{(n+k)!}{n!k!} = \frac{n+k}{nk} \frac{\Gamma(n+k)}{\Gamma(n)\Gamma(k)} = \frac{n+k}{nk} \frac{1}{B(n,k)}.
$$

By using these connections, we obtain the following bound:

$$
\begin{aligned}
\mathbb{E}[S^{\epsilon}(n, n+k)] &= \binom{n+k}{n} n! \cdot P^{\epsilon}(n, n+k) \\
&\leq n! \cdot \frac{n+k}{nk} \frac{1}{B(n,k)} \cdot \frac{e^2\, p^{-n}}{(n-1)!} B(n, k/\sqrt{p}) \\
&= e^2 \frac{(n+k)}{k} p^{-n} \frac{B(n, k/\sqrt{p})}{B(n,k)} \\
&\leq e^2 \frac{(n+k)}{k} p^{-n} = \exp(O(\log n))
\end{aligned}
$$

because $p = e^{-2\epsilon\lambda}$ and $\epsilon\lambda = O(n^{-1}\log n)$. ∎

Proposition 4. *If $\epsilon\lambda = \omega(n^{-1}\log n)$, i.e., $\epsilon\lambda \gg n^{-1}\log n$, then $\mathbb{E}[S^\epsilon(n, n+k)]$ is super-polynomially large, uniformly for all $k \geq 1$.*

Proof. Since the expected value $\mathbb{E}[S^\epsilon(n, n+k)]$ increases with $\epsilon\lambda$, and $\epsilon\lambda = \omega(n^{-1}\log n)$, it suffices to focus on $p = e^{-2\epsilon\lambda} = 1 - o(1)$ for the remainder of the proof.

Consistent with the previous analysis, we define the following variables:

$$s \equiv \sum_{i\in[n]} x_i, \quad s_j \equiv \sum_{i\neq j} x_i = s - x_j, \quad t \equiv \sum_{i\in[n]} x_i^2, \quad \text{and} \quad t_j \equiv \sum_{i\neq j} x_i^2 = t - x_j^2.$$

We restrict our attention to the set D of all non-negative $x = (x_1, x_2, \ldots, x_n)$ such that

$$\begin{aligned} &s_1(n) \leq s \leq s_2(n), \\ &s^{-1}x_i \leq \frac{2\log n}{n}, \quad i \in [n], \\ &s^{-2}t \leq \frac{3}{n}, \end{aligned} \tag{8}$$

where $s_2(n) > s_1(n)$ will be defined later. (As before, the two lower inequalities are guided by the connection between the fractions $s^{-1}x_i$ and the lengths L_i of random sub-intervals of $[0,1]$, as well as by Lemma 4.) According to the first two conditions, we have

$$x_i \leq s\frac{2\log n}{n} \leq \frac{2s_2(n)\log n}{n} \equiv \sigma_n.$$

In what follows, we choose $s_2(n)$ such that $\sigma_n \to 0$ (to be verified later). Thus, for sufficiently large n, we have $x_i \leq \sigma_n < 1$ and $D \subset [0,1]^n$. Consequently, for such large n,

$$\begin{aligned} P^\epsilon(n, n+k) \geq \tilde{P}^\epsilon(n, n+k) \equiv \int_{\mathbf{x}\in D} &\left(\prod_{j\in[n]} \left(\int_0^1 \prod_{i\neq j} (1 - px_iy_j)\, dy_j\right)\right) \\ &\times \prod_{l\in[n]} (1 - \sqrt{p}x_l)^k \, d\mathbf{x}. \end{aligned}$$

Since $1 - \alpha = \exp(-\alpha - \alpha^2(1 + O(\alpha))/2)$ as $\alpha \to 0$, for each fixed j, we have

$$\begin{aligned} \prod_{i\in[n], i\neq j} (1 - px_iy_j) &= \exp\left(-py_js_j - p^2y_j^2t_j\frac{1 + O(\sigma_n)}{2}\right) \\ &\geq \exp\left(-py_js - \frac{2p^2y_j^2t}{3}\right). \end{aligned}$$

By proceeding as in the proof of Proposition 2, and using $t/s^2 \leq 3/n$, we can bound

$$\prod_{j\in[n]}\left(\int_0^1 \prod_{i\neq j}(1-px_iy_j)\,dy_j\right) \geq e^{-4}\left(\frac{1-e^{-ps}}{ps}\right)^n.$$

Furthermore,

$$\begin{aligned}\prod_{l\in[n]}(1-\sqrt{p}x_l)^k &= \exp\left(-k\sqrt{p}s - kpt\,\frac{1+O\left(\sigma_n\right)}{2}\right)\\ &\geq \exp\left(-k\sqrt{p}s - \frac{2kps^2}{3}\frac{t}{s^2}\right)\\ &\geq \exp\left(-k\sqrt{p}s - \frac{2kps^2}{n}\right).\end{aligned}$$

Based on the previous calculations, we have

$$P^{\epsilon}(n,n+k) \geq \int_{\mathbf{x}\in D} e^{-4}\left(\frac{1-e^{-ps}}{ps}\right)^n \exp\left(-k\sqrt{p}s - \frac{2kps^2}{n}\right)\,d\mathbf{x}.$$

We then switch to new variables $(u, v_1, v_2, \ldots, v_{n-1})$ defined as follows:

$$u \equiv \sum_{i\in[n]} x_i = s \qquad \text{and} \qquad v_i \equiv s^{-1}x_i, \quad i\in[n-1].$$

Additionally, we define $v_n \equiv s^{-1}x_n = 1 - \sum_{i\in[n-1]} v_i$. The conditions in (8) then become

$$\begin{aligned}&s_1(n) \leq u \leq s_2(n),\\ &v_i \leq \frac{2\log(n)}{n}, \quad i\in[n],\\ &\sum_{i\in[n]} v_i^2 \leq \frac{3}{n}.\end{aligned} \tag{9}$$

Notably, for these new variables, the resulting region is the direct product of the range of u and the range of $\mathbf{v}$. By using equation (3), we obtain

$$\begin{aligned}\tilde{P}^{\epsilon}(n,n+k) \geq \frac{e^{-4}\,p^{-n}}{(n-1)!}&\left(\int_{s_1(n)}^{s_2(n)}(1-e^{-pu})^n\exp\left(-k\sqrt{p}u\right)\cdot\frac{\exp\left(-\frac{2kpu^2}{n}\right)}{u}du\right)\\ &\times\mathbb{P}\left(L_n^+ \leq \frac{2\log(n)}{n} \quad \text{and} \quad T_n \leq \frac{3}{n}\right),\end{aligned}$$

where $L_n^+ \equiv \max_{i\in[n]} L_i$ and $T_n \equiv \sum_{i\in[n]} L_i^2$. By Lemma 4, the probability of the event on the right-hand side tends to one as $n \to \infty$. Thus, it remains to examine

$$\int^\star \equiv \int_{s_1(n)}^{s_2(n)} \frac{\exp\left(H(u) - \frac{2kpu^2}{n}\right)}{u}\, du, \quad \text{where}$$
$$H(u) \equiv n\log(1-e^{-pu}) - k\sqrt{p}u,$$

for appropriately chosen $s_{1,2}(n)$.

The derivative of $H(u)$ is given by

$$H'(u) = \frac{np}{e^{pu}-1} - k\sqrt{p},$$

implying that the function $H(u)$ attains its maximum at $s(n) \equiv \frac{\log\left(\frac{n+k/\sqrt{p}}{k/\sqrt{p}}\right)}{p}$. In addition,

$$H''(u) = -\frac{np^2e^{pu}}{(e^{pu}-1)^2} = -\frac{np^2}{(e^{pu/2}-e^{-pu/2})^2} \geq -\frac{n}{u^2},$$

since the inequality $e^{\xi/2} - e^{-\xi/2} \geq \xi$ holds for $\xi \geq 0$.

We anticipate that the dominant contribution to the integral $\int^\star$ is obtained when we set

$$s_{1,2}(n) = s(n)(1 \mp \delta(n))$$

for some $\delta(n) \to 0$. (Hence, $\sigma_n = \frac{2s_2(n)\log n}{n} \to 0$ for $n \to \infty$, as required.) For the considered interval, $u \in [s_1(n), s_2(n)]$, we have $H''(u) \geq -\frac{n}{u^2} = -\Theta(n/s^2(n))$.

Note that

$$H(s(n)) = n\log\left(\frac{n}{n+k/\sqrt{p}}\right) - \frac{k}{\sqrt{p}}\log\left(\frac{n+k/\sqrt{p}}{k/\sqrt{p}}\right)$$

can be expressed as

$$H(s(n)) = J\left(k/\sqrt{p}\right) \quad \text{with} \quad J(z) \equiv n\log\frac{n}{n+z} + z\log\frac{z}{n+z}.$$

Furthermore, $J'(z) = \log\frac{z}{n+z}$ and $J''(z) = \frac{n}{z(n+z)} > 0$. By using the tangent line inequality $J(z) \geq J(k) + J'(k)(z-k)$ for $z = k/\sqrt{p}$, we get

$$H(s(n)) \geq n\log\frac{n}{n+k} + k\log\frac{k}{n+k} - \frac{1-p}{p+\sqrt{p}}k\log\frac{n+k}{k}.$$

Thus, by expanding $H(u)$ around its maximum $u = s(n)$

$$H(u) = H(s(n)) + \frac{H''(\xi)}{2}(u - s(n))^2, \quad \text{where } \xi \text{ is between } u \text{ and } s(n),$$

we can conclude that $\int^*$ is at least of order

$$\frac{\exp(H(s(n)) - 3kps^2(n)/n)}{s(n)} \int_{-\delta(n)s(n)}^{\delta(n)s(n)} \exp\left(-z^2\Theta(n/s^2(n)\right)\, dz.$$

By choosing $\delta(n) \to 0$ such that $\delta(n)\sqrt{n} \to \infty$, the last integral equals

$$\int_{-\delta(n)s(n)}^{\delta(n)s(n)} \exp\left(-z^2\Theta(n/s^2(n)\right)\, dz = \frac{s(n)}{\sqrt{n}} \int_{-\delta(n)\sqrt{n}}^{\delta(n)\sqrt{n}} e^{-\Theta\left(\eta^2\right)}\, d\eta = \Theta\left(\frac{s(n)}{\sqrt{n}}\right).$$

Therefore, the integral $\int^*$ is at least of order

$$\frac{\exp(H(s(n)) - 3kps^2(n)/n)}{\Theta(\sqrt{n})} = \frac{n^n k^k}{(n+k)^{n+k}\,\Theta(\sqrt{n})} \times \exp\left(-\frac{1-p}{p+\sqrt{p}} k \log\frac{n+k}{k} - 3kps^2(n)/n\right).$$

Hence, $\tilde{P}^\epsilon(n, n+k)$ is at least of order

$$\frac{p^{-n}}{(n-1)!}\frac{n^n k^k}{(n+k)^{n+k}\sqrt{n}} \exp\left(-\frac{1-p}{p+\sqrt{p}} k \log\frac{n+k}{k} - 3kps^2(n)/n\right).$$

Using

$$\mathbb{E}[S^\epsilon(n, n+k)] \geq \binom{n+k}{n} n! \cdot \tilde{P}^\epsilon(n, n+k) \quad \text{and} \quad \binom{n+k}{n} = \Omega\left(\frac{1}{\sqrt{n}}\,\frac{(n+k)^{n+k}}{n^n k^k}\right),$$

we obtain that $\mathbb{E}[S^\epsilon(n, n+k)]$ is at least of order

$$\exp\left(n\log(1/p) - \frac{1-p}{p+\sqrt{p}} k \log\frac{n+k}{k} - 3kps^2(n)/n\right).$$

In this last expression, by the definition of $s(n)$, we have

$$kps^2(n)/n = \frac{1}{\sqrt{p}} \cdot \frac{\log^2(1+x)}{x}\bigg|_{x=\frac{n}{k/\sqrt{p}}} = O(1/\sqrt{p}).$$

In addition,

$$\frac{k}{n}\log\frac{n+k}{k} = \frac{\log(1+x)}{x}\bigg|_{x=\frac{n}{k}} \leq 1 \quad \text{and} \quad \log(1/p) > 1-p.$$

To sum up, $\mathbb{E}[S^\epsilon(n, n+k)]$ is at least of order

$$\begin{aligned}\exp\left(n\left(\log(1/p) - \frac{1-p}{p+\sqrt{p}}\right) + O(1/\sqrt{p})\right) &\geq \exp\left(n(1-p)\frac{p+\sqrt{p}-1}{p+\sqrt{p}} + O(1/\sqrt{p})\right) \\ &= \exp(\omega(\log n))\end{aligned}$$

since $p = e^{-2\epsilon\lambda}$, $\epsilon\lambda = o(1)$, and $\epsilon\lambda = \omega(n^{-1}\log n)$. ∎

Remark. We can employ even tighter arguments to derive sharper bounds. By refining our analysis, it appears feasible to obtain an asymptotic expression for the expectation $\mathbb{E}[S^\epsilon(n, n+k)]$, particularly when $\epsilon\lambda = o(1)$; this represents one direction of our current research. To provide a glimpse of what such an expression might look like, without delving into technical details, note that the integral $\int^\star$, as defined in the proof of Proposition 4, can also be written as

$$\int^\star = \frac{\exp(-2kps^2(n)/n)}{ps(n)} \int_{z_1}^{z_2} z^{\frac{k}{\sqrt{p}}-1}(1-z)^n\, dz, \quad \text{where} \quad z_{1,2} \equiv \left(\frac{k/\sqrt{p}}{n+k/\sqrt{p}}\right)^{1\pm\delta(n)}.$$

The last integral closely resembles the beta function $B(n+1, k/\sqrt{p})$ and, in fact, may be well-approximated by this expression for appropriately chosen $\delta(n) \to 0$. This makes the lower bound conceptually closer to the upper bound examined in the proof of Proposition 3, also reducing the gap between the two. Further analysis could narrow this gap even more.

Appendix C. Additional Integral Formulas

A notable advantage of our model is its ability to yield a variety of tractable integral formulas for key statistics of interest. These formulas appear particularly amenable to further examination, including asymptotic analysis. As an example, in this section, we derive integral formulas that allow the study of the average rank of workers, or firms, as ranked by their respective partners. The asymptotic properties of these and other related statistics are key areas of our current focus.

For simplicity of exposition, consider first the balanced market with $k = 0$. Using the transformation of random markets introduced in Appendix A, we consider $U^f_{i,j} = \frac{1}{\lambda}\log\frac{1}{X_{i,j}}$ and $U^w_{i,j} = \frac{1}{\lambda}\log\frac{1}{Y_{i,j}}$, where $\{X_{i,j}, Y_{i,j}\}_{i,j\in[n]}$ are $2n^2$ i.i.d. random variables uniformly distributed on $[0,1]$. Focus on the diagonal matching μ and let

$$\begin{aligned} TR(\mu) &\equiv n + \sum_{i=1}^{n} \left| \left\{ j \leq n : \, U^f_{i,j} > U^f_{i,i} \right\} \right| \\ &= n + \sum_{i=1}^{n} \left| \left\{ j \leq n : \, \frac{1}{\lambda}\log\frac{1}{X_{i,j}} > \frac{1}{\lambda}\log\frac{1}{X_{i,i}} \right\} \right| \\ &= n + \sum_{i=1}^{n} \left| \{ j \leq n : \, X_{i,j} < X_{i,i} \} \right| \end{aligned}$$

denote the total rank of workers as ranked by firms, which counts the number of workers whom firms find as good as their respective partners under μ. (*Remark.* Alternatively, we could count workers whom firms ϵ-prefer to their partners.)

Define

$$P^{\epsilon}_{a}(n,n) \equiv \mathbb{P}(\mu \text{ is } \epsilon\text{-stable and } TR(\mu) = a), \quad a \in [n, n^2],$$

as the probability that the diagonal matching μ is ϵ-stable and that the total rank of workers under μ is equal to a. By symmetry, the same probability applies to any complete matching, not necessarily the diagonal one.

Lemma 5. *Consider a random market with n firms and n workers for arbitrary costs $\epsilon \geq 0$. Then,*

$$P^{\epsilon}_{a}(n,n) = \int\limits_{\mathbf{x},\mathbf{y}\in[0,1]^n} [\alpha^{a-n}] \prod_{i\neq j} \left(1 - x_i(1 - \alpha + e^{-2\epsilon\lambda} y_j \alpha)\right) \, d\mathbf{x} d\mathbf{y},$$

where $[\alpha^{a-n}]$ denotes the coefficient of α^{a-n}. In particular, this probability is a function of the relative switching cost $\epsilon\lambda = \epsilon / \mathbb{E}[Exp\,(\lambda)]$.

Naturally, by dropping $[\alpha^{a-n}]$ and setting $\alpha = 1$, we obtain the integral formula for $P^{\epsilon}(n,n)$, derived in Lemma 2.

Proof of Lemma 5. We first determine $P_a^{\epsilon}(\mu|\circ)$, the conditional probability of the desired event $\{\mu \text{ is } \epsilon\text{-stable and } TR(\mu) = a\}$ given $X_{i,i} = x_i$ and $Y_{j,j} = y_j$, where $i, j \in [n]$. Using generating functions, we can rewrite this as

$$P_a^{\epsilon}(\mu|\circ) = [\alpha^{a-n}]\ \mathbb{E}[\mathbb{1}(\mu \text{ is } \epsilon\text{-stable})\, \alpha^{TR(\mu)-n}|\circ].$$

We evaluate this expression by using the following "marking" procedure. Fix $\alpha \in (0,1)$. Scan all pairs $(i,j) \in [n]^2$ and, whenever $\frac{1}{\lambda}\log\frac{1}{X_{i,j}} > \frac{1}{\lambda}\log\frac{1}{x_i}$ (i.e., $X_{i,j} < X_{i,i} = x_i$) holds, mark the pair (i,j) with probability α, independently of all other pairs.

Then,

$$\mathbb{E}[\mathbb{1}(\mu \text{ is } \epsilon\text{-stable})\, \alpha^{TR(\mu)-n}|\circ] = \mathbb{P}(\mathcal{B}|\circ),$$

where $\mathcal{B}$ denotes the event that μ is ϵ-stable and all pairs (i,j) with $X_{i,j} < X_{i,i} = x_i$ are marked. Note that

$$\mathcal{B} = \bigcap_{i \neq j} \mathcal{B}_{i,j},$$

with

$$\begin{aligned}\mathcal{B}_{i,j} =& \left\{X_{i,j} > x_i\right\} \quad \cup \\ & \left\{X_{i,j} < x_i,\ \frac{1}{\lambda}\log\frac{1}{X_{i,j}} < \frac{1}{\lambda}\log\frac{1}{x_i} + \epsilon,\ (i,j) \text{ is marked}\right\} \quad \cup \\ & \left\{\frac{1}{\lambda}\log\frac{1}{X_{i,j}} > \frac{1}{\lambda}\log\frac{1}{x_i} + \epsilon,\ \frac{1}{\lambda}\log\frac{1}{Y_{i,j}} < \frac{1}{\lambda}\log\frac{1}{y_j} + \epsilon,\ (i,j) \text{ is marked}\right\}.\end{aligned}$$

That is,

$$\begin{aligned}\mathcal{B}_{i,j} =& \left\{X_{i,j} > x_i\right\} \quad \cup \\ & \left\{e^{-\epsilon\lambda}x_i < X_{i,j} < x_i,\ (i,j) \text{ is marked}\right\} \quad \cup \\ & \left\{X_{i,j} < e^{-\epsilon\lambda}x_i,\ Y_{i,j} > e^{-\epsilon\lambda}y_j,\ (i,j) \text{ is marked}\right\}.\end{aligned}$$

The events $\mathcal{B}_{i,j}$, $i \neq j$, are conditionally independent.

Furthermore,

$$\begin{aligned}\mathbb{P}(\mathcal{B}_{i,j}|\circ) &= (1-x_i) + x_i(1-e^{-\epsilon\lambda})\alpha + e^{-\epsilon\lambda}x_i(1-e^{-\epsilon\lambda}y_j)\alpha \\ &= 1 - x_i + x_i\alpha - e^{-\epsilon\lambda}x_i\alpha + e^{-\epsilon\lambda}x_i\alpha - e^{-2\epsilon\lambda}x_iy_j\alpha \\ &= 1 - x_i(1-\alpha+e^{-2\epsilon\lambda}y_j\alpha).\end{aligned}$$

By collecting all the pieces and integrating over $\mathbf{x}, \mathbf{y} \in [0,1]^n$, we obtain the desired integral formula. ∎

Lemma 6 provides integral formulas for general markets; its proof is more involved and omitted for brevity. Specifically, $P_a^\epsilon(n, n+k)$ denotes the probability that a matching—with all firms matched—is stable and the firms' total rank of their paired workers is a. Similarly, $P_b^\epsilon(n, n+k)$ represents the probability that a matching is stable and the matched workers' total rank of their paired firms is b. Finally, $P_{a,b}^\epsilon(n, n+k)$ is the probability that a matching is stable, where the firms' total rank is a and the matched workers' total rank is b.

Lemma 6. *Consider a random market with n firms and $(n+k)$ workers for arbitrary imbalance $k \geq 0$ and costs $\epsilon \geq 0$. Then,*

$$P_a^\epsilon(n,n+k) = \int_{\mathbf{x},\mathbf{y}\in[0,1]^n} [\alpha^{a-n}] \prod_{i\neq j} \left(1 - x_i(1-\alpha+e^{-2\epsilon\lambda}y_j\alpha)\right) \prod_{l\in[n]} (1 + x_l((1-e^{-\epsilon\lambda})\alpha - 1))^k \, d\mathbf{x}d\mathbf{y},$$

$$P_b^\epsilon(n,n+k) = \int_{\mathbf{x},\mathbf{y}\in[0,1]^n} [\beta^{b-n}] \prod_{i\neq j} \left(1 - y_j(1-\beta+e^{-2\epsilon\lambda}x_i\beta)\right) \prod_{l\in[n]} (1 - e^{-\epsilon\lambda}x_l)^k \, d\mathbf{x}d\mathbf{y}, \quad \textit{and}$$

$$\begin{aligned}P_{a,b}^\epsilon(n,n+k) = \int_{\mathbf{x},\mathbf{y}\in[0,1]^n} [\alpha^{a-n}\beta^{b-n}] &\prod_{i\neq j} \left(1 + x_i(\alpha-1) + y_j(\beta-1) + x_iy_j(1-\alpha-\beta+(1-e^{-2\epsilon\lambda})\alpha\beta)\right) \\ &\times \prod_{l\in[n]} (1 + x_l((1-e^{-\epsilon\lambda})\alpha - 1))^k \, d\mathbf{x}d\mathbf{y}.\end{aligned}$$

Lemma 7 is an analogue of Lemma 6 but relies on modified ranks. For each matched agent, a modified rank counts the number of agents from the other market side whom the agent ϵ-prefers over their current partner—plus one to include their current partner in this count by convention. In comparison, a standard rank for each matched agent counts the number of agents whom the agent views as at least as good as their current partner. Let $\mathcal{P}_a^\epsilon(n, n+k)$, $\mathcal{P}_b^\epsilon(n, n+k)$, and $\mathcal{P}_{a,b}^\epsilon(n, n+k)$ denote the counterparts of $P_a^\epsilon(n, n+k)$, $P_b^\epsilon(n, n+k)$, and $P_{a,b}^\epsilon(n, n+k)$, respectively, under this modified ranking approach.

Lemma 7. *Consider a random market with n firms and $(n+k)$ workers for arbitrary imbalance $k \geq 0$ and costs $\epsilon \geq 0$. Then,*

$$\mathcal{P}_a^\epsilon(n, n+k) = \int_{\mathbf{x},\mathbf{y}\in[0,1]^n} [\alpha^{a-n}] \prod_{i\neq j} \left(1 - e^{-\epsilon\lambda} x_i (1 - \alpha + e^{-\epsilon\lambda} y_j \alpha)\right) \prod_{l\in[n]} (1 - e^{-\epsilon\lambda} x_l)^k \, d\mathbf{x} d\mathbf{y},$$

$$\mathcal{P}_b^\epsilon(n, n+k) = \int_{\mathbf{x},\mathbf{y}\in[0,1]^n} [\beta^{b-n}] \prod_{i\neq j} \left(1 - e^{-\epsilon\lambda} y_j (1 - \beta + e^{-\epsilon\lambda} x_i \beta)\right) \prod_{l\in[n]} (1 - e^{-\epsilon\lambda} x_l)^k \, d\mathbf{x} d\mathbf{y}, \quad \textit{and}$$

$$\mathcal{P}_{a,b}^\epsilon(n, n+k) = \int_{\mathbf{x},\mathbf{y}\in[0,1]^n} [\alpha^{a-n}\beta^{b-n}] \prod_{i\neq j} \left(1 + e^{-\epsilon\lambda} x_i(\alpha - 1) + e^{-\epsilon\lambda} y_j(\beta - 1) + e^{-2\epsilon\lambda} x_i y_j (1 - \alpha - \beta)\right) \times \prod_{l\in[n]} (1 - e^{-\epsilon\lambda} x_l)^k \, d\mathbf{x} d\mathbf{y}.$$

Appendix D. Proof of Lemma 3 and Non-Integrality

In this section, we prove the results stated in Section 4. We first provide the proof of Lemma 3. Then, Example 2 presents a matching market where a non-integer point of the ϵ-stable polytope cannot be expressed as a linear combination of ϵ-stable matchings, demonstrating that the polytope is not integral.

Proof of Lemma 3. Note that the condition (2)

$$x_{i,j} + \sum_{k\neq j:\, u_{i,k}^f + \epsilon \geq u_{i,j}^f} x_{i,k} + \sum_{k\neq i:\, u_{k,j}^w + \epsilon \geq u_{i,j}^w} x_{k,j} \geq 1 \quad \text{for each } (f_i, w_j) \in \mathcal{F} \times \mathcal{W},$$

is violated if and only if, for some $(f_i, w_j) \in \mathcal{F} \times \mathcal{W}$, the following three equalities hold:

$$\begin{aligned}
x_{i,j} = 0 \quad &\iff \quad f_i \text{ and } w_j \text{ are not matched,} \\
\sum_{k\neq j:\, u_{i,k}^f + \epsilon \geq u_{i,j}^f} x_{i,k} = 0 \quad &\iff \quad f_i \text{ is single or matched to } w_k \text{ with } u_{i,k}^f + \epsilon < u_{i,j}^f, \\
\sum_{k\neq i:\, u_{k,j}^w + \epsilon \geq u_{i,j}^w} x_{k,j} = 0 \quad &\iff \quad w_j \text{ is single or matched to } f_k \text{ with } u_{k,j}^w + \epsilon < u_{i,j}^w.
\end{aligned}$$

These conditions, in turn, hold if and only if the pair (f_i, w_j) is an ϵ-blocking pair. ∎

Example 2. Consider the following market with three firms and three workers:[27]

$$\begin{array}{c} \\ f_1 \\ f_2 \\ f_3 \end{array} \begin{array}{c} \begin{array}{ccc} w_1 & w_2 & w_3 \end{array} \\ \begin{pmatrix} 6,1 & 5,6 & 4,6 \\ 6,6 & 10,5 & 2,5 \\ 6,5 & 5,4 & 4,4 \end{pmatrix} \end{array}.$$

It is straightforward to verify that, for switching costs $\epsilon = 3$, the vector

$$x = \begin{pmatrix} 1/2 & 0 & 1/2 \\ 0 & 1/2 & 1/2 \\ 1/2 & 1/2 & 0 \end{pmatrix}$$

belongs to the ϵ-stable polytope; namely, it satisfies (1) and (2).

Furthermore, this vector x is uniquely decomposed into (integer) matchings:

$$x = \frac{1}{2}\, y + \frac{1}{2}\, z,$$

where

$$y = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad \text{and} \quad z = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}.$$

Nonetheless, in this unique decomposition, matching z is not ϵ-stable. Indeed, firm f_2 and worker w_1 form an ϵ-blocking pair. Thus, there has to be a non-integer extreme point, implying that the ϵ-stable polytope is not integral. △

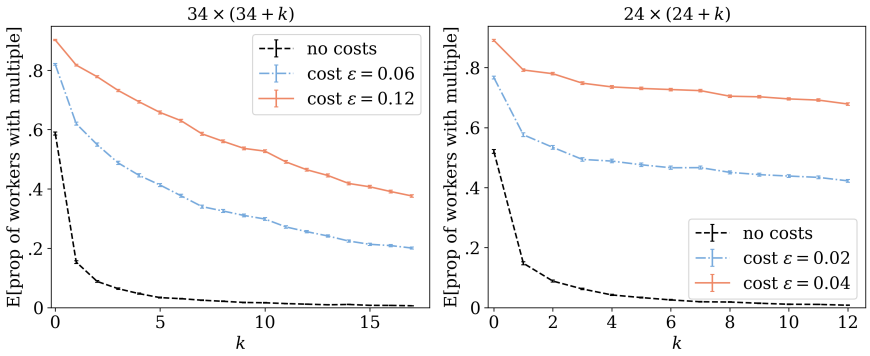
Appendix E. Additional Figures

In this section, we present additional computational results that complement those in Section 4, illustrating how switching costs impact variability in matching outcomes and can provide welfare gains to matched workers.

Figure 6 is the worker counterpart to Panels (3a) and (4a) in the main text and displays the proportion of workers that have different partners across ϵ-stable matchings. In contrast to a frictionless scenario, even with small switching costs, many workers have multiple partners—this accounts for workers who are matched in one ϵ-stable matching but

[27] This example is inspired by Example 2 in Aziz and Klaus (2019).

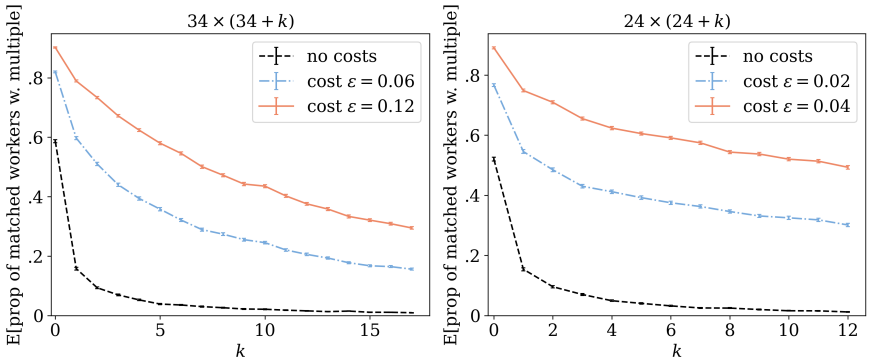
unmatched in another, with being unmatched treated as a distinct partner.



(a) Exponential (b) Uniform

Figure 6: Proportion of workers with multiple partners

Figure 7 is similar to Figure 6 but focuses exclusively on workers who are consistently matched across all ϵ-stable matchings. As shown, a sizable fraction of these consistently matched workers have multiple partners—reinforcing the insights from Panels (3c) and (4c) in the main text, which demonstrate that these partners can be substantially different.

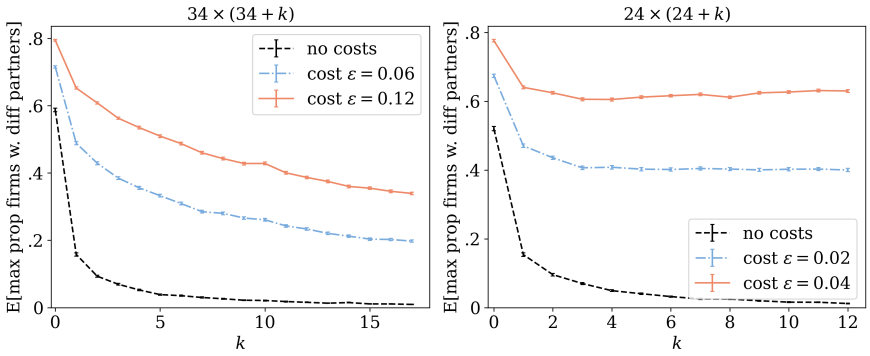


(a) Exponential (b) Uniform

Figure 7: Proportion of always matched workers with multiple partners

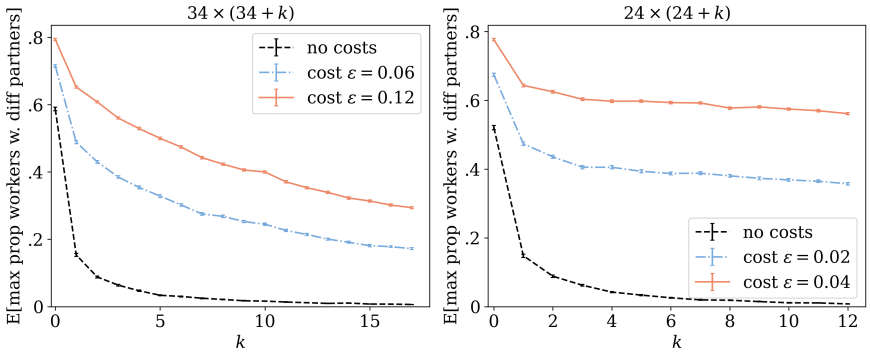
Figures 8 and 9 show that ϵ-stable matchings can involve a sizable proportion of mismatched agents compared to an almost unique stable matching in a scenario with no switching costs. As a benchmark for this almost unique stable matching, we focus on the stable matching produced by the firm-proposing Deferred Acceptance (F-DA) algorithm (Gale and Shapley, 1962). These figures display the maximal proportions of firms and workers, respectively, that have different partners in some ϵ-stable matching compared to the benchmark,

almost unique stable matching.[28]



(a) Exponential (b) Uniform

Figure 8: Maximal proportion of firms with different partners compared to those in the F-DA outcome

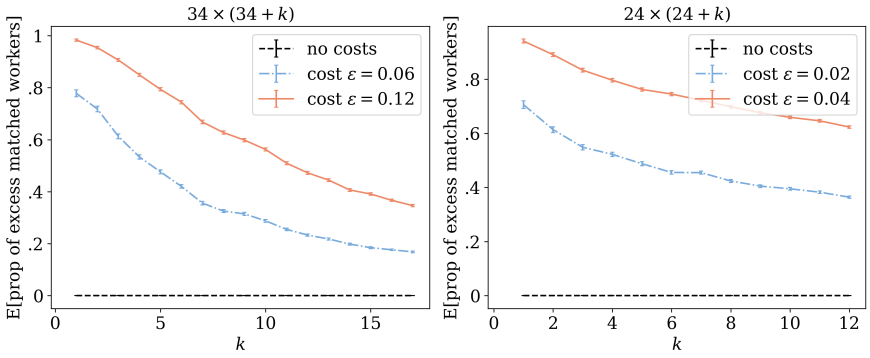


(a) Exponential (b) Uniform

Figure 9: Maximal proportion of workers with different partners compared to those in the F-DA outcome

Figure 10 complements Panels (3d) and (4d), showing the proportion of workers who can be matched in some ϵ-stable matching among workers who would remain (consistently) unmatched in the absence of switching costs. Naturally, this proportion is zero in a frictionless scenario, where the Rural Hospital Theorem guarantees a consistent set of matched agents across all stable matchings. Nonetheless, as can be seen, these proportions may be sizable even with small switching costs, especially when workers are more replaceable.

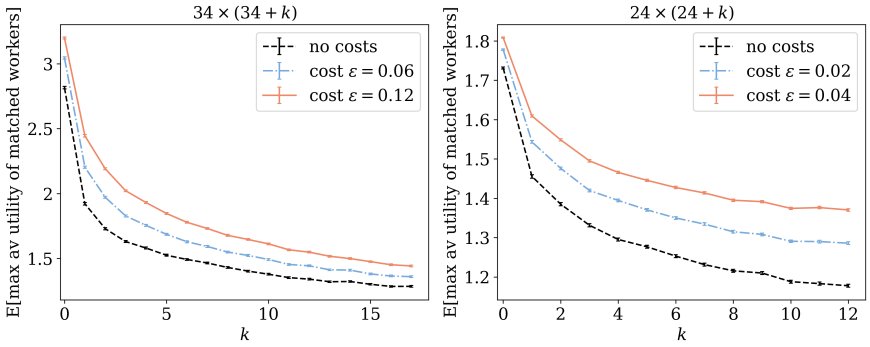
[28] Additionally, Figure 8, along with Panels (3a) and (4a) in the main text, provides bounds on the maximal proportion of firms that have different partners across two ϵ-stable matchings. Similarly, Figures 9 and 6 offer analogous bounds for workers.



(a) Exponential (b) Uniform

Figure 10: Excess fraction of potentially matched workers

Figure 11 shows that even small switching costs can lead to significant average welfare gains—particularly compared to the costs—for matched workers. Switching costs create room for selecting which workers get matched and determining how they are matched, resulting in considerably higher utilities for matched workers compared to a scenario without costs. The selection channel is especially relevant for the uniform distribution, where workers are more replaceable, and this replaceability becomes more relevant with greater imbalances.



(a) Exponential (b) Uniform

Figure 11: Maximal average utility of matched workers

References

Abdulkadiroğlu, A. and T. Sönmez (2003): "School Choice: A Mechanism Design Approach," *American Economic Review*, 93, 729–747.

Abraham, D. J., P. Biró, and D. F. Manlove (2005): ""Almost Stable" Matchings in the Roommates Problem," in *International Workshop on Approximation and Online Algorithms*, Springer, 1–14.

Agarwal, N. and P. Somaini (2023): "Empirical Models of Non-Transferable Utility Matching," *Online and Matching-Based Market Design*.

Alon, N., K. Rudov, and L. Yariv (2024): "Dominance Solvability in Random Games," *mimeo*.

Anshelevich, E., S. Das, and Y. Naamad (2013): "Anarchy, Stability, and Utopia: Creating Better Matchings," *Autonomous Agents and Multi-Agent Systems*, 26, 120–140.

Artemov, G., Y.-K. Che, and Y. He (2023): "Stable Matching with Mistaken Agents," *Journal of Political Economy Microeconomics*, 1, 270–320.

Ashlagi, I., M. Braverman, Y. Kanoria, and P. Shi (2020): "Clearing Matching Markets Efficiently: Informative Signals and Match Recommendations," *Management Science*, 66, 2163–2193.

Ashlagi, I., Y. Kanoria, and J. D. Leshno (2017): "Unbalanced Random Matching Markets: The Stark Effect of Competition," *Journal of Political Economy*, 125, 69–98.

Azar, J. and I. Marinescu (2024): "Monopsony Power in the Labor Market: From Theory to Policy," *Annual Review of Economics*, 16, 491–518.

Aziz, H., P. Biró, S. Gaspers, R. de Haan, N. Mattei, and B. Rastegari (2020): "Stable Matching with Uncertain Linear Preferences," *Algorithmica*, 82, 1410–1433.

Aziz, H. and B. Klaus (2019): "Random Matching under Priorities: Stability and No Envy Concepts," *Social Choice and Welfare*, 53, 213–259.

Balinski, M. and T. Sönmez (1999): "A Tale of Two Mechanisms: Student Placement," *Journal of Economic theory*, 84, 73–94.

Biro, P., A. Hassidim, A. Romm, R. I. Shorrer, and S. Sovago (2022): "The Large Core of College Admission Markets: Theory and Evidence," in *Proceedings of the 23rd ACM Conference on Economics and Computation*, 958–959.

Budish, E. and E. Cantillon (2012): "The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," *American Economic Review*, 102, 2237–2271.

Bursztyn, L. and R. Jensen (2017): "Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure," *Annual Review of Economics*, 9, 131–153.

Caragiannis, I., A. Filos-Ratsikas, P. Kanellopoulos, and R. Vaish (2021): "Stable Fractional Matchings," *Artificial Intelligence*, 295, 103416–103416.

Che, Y.-K. and O. Tercieux (2019): "Efficiency and Stability in Large Matching Markets," *Journal of Political Economy*, 127, 2301–2342.

Chen, J., P. Skowron, and M. Sorge (2021): "Matchings under Preferences: Strength of Stability and Tradeoffs," *ACM Transactions on Economics and Computation*, 9, 1–55.

Chiappori, P.-A. and B. Salanié (2016): "The Econometrics of Matching Models," *Journal of Economic*

Literature, 54, 832–861.

Chu, C. S., P. Leslie, and A. Sorensen (2011): "Bundle-Size Pricing as an Approximation to Mixed Bundling," *The American Economic Review*, 263–303.

Combe, J., U. Dur, O. Tercieux, C. Terrier, and U. M. Ünver (2022): "Market Design for Distributional Objectives in (Re)assignment: An Application to Improve the Distribution of Teachers in Schools," *mimeo*.

Crawford, G. S. and A. Yurukoglu (2012): "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 102, 643–685.

Diss, M. and V. Merlin (2021): "Evaluating Voting Systems with Probability Models," *Studies in Choice and Welfare. Springer International Publishing*.

Doğan, B. and L. Ehlers (2021): "Minimally Unstable Pareto Improvements over Deferred Acceptance," *Theoretical Economics*, 16, 1249–1279.

Donna, J. D. and J.-A. Espín-Sánchez (2018): "Complements and Substitutes in Sequential Auctions: The Case of Water Auctions," *The RAND Journal of Economics*, 49, 87–127.

Echenique, F., R. Gonzalez, A. J. Wilson, and L. Yariv (2022): "Top of the Batch: Interviews and the Match," *American Economic Review: Insights*, 4, 223–38.

Echenique, F., N. Immorlica, and V. V. Vazirani (2023): *Online and Matching-Based Market Design*, Cambridge University Press.

Echenique, F., A. Robinson-Cortés, and L. Yariv (2024): "An Experimental Study of Decentralized Matching," *mimeo*.

Erdil, A. and H. Ergin (2017): "Two-Sided Matching with Indifferences," *Journal of Economic Theory*, 171, 268–292.

Feller, W. (1971): *An Introduction to Probability Theory and Its Applications*, Wiley, New York, 2nd ed.

Fernandez, M. A., K. Rudov, and L. Yariv (2022): "Centralized Matching with Incomplete Information," *American Economic Review: Insights*, 4, 18–33.

Fox, J. T. (2009): "Structural Empirical Work Using Matching Models," *New Palgrave Dictionary of Economics. Online edition*.

Gale, D. and L. S. Shapley (1962): "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 69, 9–15.

Gangam, R. R., T. Mai, N. Raju, and V. V. Vazirani (2022): "A Structural and Algorithmic Study of Stable Matching Lattices of "Nearby" Instances, with Applications," *arXiv e-prints*, arXiv–1804.

Genc, B., M. Siala, G. Simonin, and B. O'Sullivan (2019): "Complexity Study for the Robust Stable Marriage Problem," *Theoretical Computer Science*, 775, 76–92.

Hall, R. E. and A. B. Krueger (2012): "Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-Job Search," *American Economic Journal: Macroeconomics*, 4, 56–67.

Hartline, J. D. and T. Roughgarden (2014): "Optimal Platform Design," *arXiv preprint arXiv:1412.8518*.

Hassidim, A., D. Marciano, A. Romm, and R. I. Shorrer (2017): "The Mechanism Is Truthful, Why

Aren't You?" *American Economic Review*, 107, 220–224.

Hoffman, C., A. Levy, and E. Mossel (2023): "Stable Matchings with Correlated Preferences," *arXiv preprint arXiv:2312.14813*.

Immorlica, N. and M. Mahdian (2005): "Marriage, Honesty, and Stability," in *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, 53–62.

Irving, R. W. and P. Leather (1986): "The Complexity of Counting Stable Marriages," *SIAM Journal on Computing*, 15, 655–667.

Kahn, L. B. (2010): "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy," *Labour economics*, 17, 303–316.

Kanoria, Y., S. Min, and P. Qian (2023): "The Competition for Partners in Matching Markets," *Minor Revision in Management Science*.

Kanoria, Y., D. Saban, and J. Sethuraman (2018): "Convergence of the Core in Assignment Markets," *Operations Research*, 66, 620–636.

Kim, D.-H. (2015): "Flexible Bayesian Analysis of First Price Auctions Using a Simulated Likelihood," *Quantitative Economics*, 6, 429–461.

Knuth, D. E. (1976): *Marriages Stables*, Les Presses de l'université de Montréal.

Kojima, F. and P. A. Pathak (2009): "Incentives and Stability in Large Two-Sided Matching Markets," *American Economic Review*, 99, 608–27.

Lauermann, S. (2013): "Dynamic Matching and Bargaining Games: A General Approach," *American Economic Review*, 103, 663–689.

Lee, S. (2017): "Incentive Compatibility of Large Centralized Matching Markets," *The Review of Economic Studies*, 84, 444–463.

Lee, S. and L. Yariv (2018): "On the Efficiency of Stable Matchings in Large Markets," *mimeo*.

Leshno, J. (2022): "Stable Matching with Peer-Dependent Preferences in Large Markets: Existence and Cutoff Characterization," *mimeo*.

Leshno, J. D. (2023): "Large Matching Markets," *Online and Matching-Based Market Design*.

McLennan, A. (2005): "The Expected Number of Nash Equilibria of a Normal Form Game," *Econometrica*, 73, 141–174.

McVitie, D. G. and L. B. Wilson (1970): "Stable Marriage Assignment for Unequal Sets," *BIT*, 10, 295–309.

——— (1971): "The Stable Marriage Problem," *Communications of the ACM*, 14, 486–490.

Nisan, N., T. Roughgarden, E. Tardos, and V. V. Vazirani (2007): *Algorithmic Game Theory*, New York, NY, USA: Cambridge University Press.

Oyer, P. (2006): "Initial Labor Market Conditions and Long-Term Outcomes for Economists," *Journal of Economic Perspectives*, 20, 143–160.

Pęski, M. (2022): "Tractable Model of Dynamic Many-to-Many Matching," *American Economic Journal: Microeconomics*, 14, 1–43.

Pittel, B. (1989): "The Average Number of Stable Matchings," *SIAM Journal on Discrete Mathematics*, 2, 530–549.

——— (1992): "On Likely Solutions of a Stable Marriage Problem," *The Annals of Applied Probability*, 2, 358–401.

——— (2019): "On Likely Solutions of the Stable Matching Problem with Unequal Numbers of Men and Women," *Mathematics of Operations Research*, 44, 122–146.

Rheingans-Yoo, R. (2024): "Large Random Matching Markets with Localized Preference Structures Can Exhibit Large Cores," *Games and Economic Behavior*, 144, 71–83.

Roth, A. E. and E. Peranson (1999): "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, 89, 748–780.

Roth, A. E., U. G. Rothblum, and J. H. Vande Vate (1993): "Stable Matchings, Optimal Assignments, and Linear Programming," *Mathematics of Operations Research*, 18, 803–828.

Roth, A. E. and M. A. O. Sotomayor (1992): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Econometric Society Monographs, Cambridge: Cambridge University Press.

Rudov, K. (2024): "Fragile Stable Matchings," *mimeo*.

Teo, C.-P. and J. Sethuraman (1998): "The Geometry of Fractional Stable Matchings and Its Applications," *Mathematics of Operations Research*, 23, 874–891.

Vande Vate, J. H. (1989): "Linear Programming Brings Marital Bliss," *Operations Research Letters*, 8, 147–153.

Wang, X., W. Zhang, S. Sheng, R. Sharma, and M. Kamran (2022): "Residency Attrition and Associated Characteristics, a 10-Year Cross Specialty Comparative Study," *Journal of Brain and Neurological Disorders*, 5.

Zhou, J. (2021): "Mixed Bundling in Oligopoly Markets," *Journal of Economic Theory*, 194, 105257.