A Perfectly Robust Approach to Multiperiod Matching Problems

Faculty Research Working Paper Series

Maciej H. Kotowski
Harvard Kennedy School

May 2019 (Updated February 2020)

RWP19-016

A Perfectly Robust Approach to Multiperiod Matching Problems

Maciej H. Kotowski[†]

February 26, 2020

Abstract

Many two-sided matching situations involve multiperiod interaction. Traditional cooperative solutions, such as pairwise stability or the core, often identify unintuitive outcomes (or are empty) when applied to such markets. As an alternative, this study proposes the criterion of perfect α-stability. An outcome is perfect α-stable if no coalition prefers an alternative assignment in any period that is superior for all plausible market continuations. The solution posits that agents have foresight, but cautiously evaluate possible future outcomes. A perfect α-stable matching exists, even when assignments are intertemporal complements. The perfect α-core, a stronger solution, is nonempty under standard regularity conditions, such as history-independence. Our analysis extends to markets with arrivals and departures, transfers, and many-to-one assignments.

Keywords: Matching, Two-sided Market, Stability, Core
JEL: C78, C71, D47

I thank Ivan Balbuzanov, David Delacrétaz, Sangram Kadam, Morimitsu Kurino, Ce Liu, Ran Shorrer, Alex Teytelboym, Norov Tumennasan, and Richard Zeckhauser for helpful comments and discussions concerning this research. I thank seminar audiences at McGill University and Monash University and participants of the 2016 Econometric Society North American Winter Meeting (San Francisco), the 2018 Meeting of the Society for Social Choice and Welfare (Seoul), the 2019 International Conference on Game Theory (Stony Brook), and the 2019 Midwest Economic Theory Conference (WUSTL) for constructive feedback. This paper incorporates and supersedes results and ideas originally communicated in the 2015 working paper "A Note on Stability in One-to-One, Multi-period Matching Markets" (HKS Faculty Research Working Paper RWP15-042).

[†]John F. Kennedy School of Government, Harvard University, 79 JFK Street, Cambridge MA 02138, United States. E-mail: <maciej_kotowski@hks.harvard.edu>

1 Introduction

In many markets, agents are partitioned into two groups—men and women, firms and workers, students and schools—and must match together to realize benefits. A common feature is that these markets' participants interact over a period of time. Matches may be fleeting or may last several periods; a few are irrevocable, but many permit revision and change.

This paper extends Gale and Shapley's (1962) model of a two-sided matching market by allowing agents to interact over multiple periods, possibly changing assignments with time. Though a natural extension of a classic setting, there is no unequivocal analogue to Gale and Shapley's stability concept to guide analysis. This paper's contribution is a family of solutions for the analysis of multiperiod two-sided matching problems with limited commitment. Models with this structure been proposed to study student-school assignment with multi-child households (Dur, 2012), teacher assignment (Pereyra, 2013), the assignment of children to daycares (Kennes et al., 2014, 2019), and the matching of medical residents to hospitals (Kadam and Kotowski, 2018a; Liu, 2018). The changing nature of matches is a key feature of the foster care and adoption system in the United States (MacDonald, 2019). Further applications include labor markets and ride-sharing platforms. Market mechanisms that coordinate agents on stable outcomes are more durable (Roth, 2002). Therefore, identifying appropriate formulations of stability for such situations has immense practical relevance.

As context for our study, it is helpful to see why a benchmark approach to dynamic two-sided matching markets is a topic of continuing debate, with many recent proposals. (We discuss the literature below.) A multiperiod matching problem straddles two conceptually-distinct structures that have been traditionally examined with tools from different paradigms. Reconciling the resulting mishmash of intuitions is not straightforward.

First, ignoring the market's time dimension, any plausible solution must resolve the standard two-sided matching problem. Here, cooperative solutions, such as stability or the core, have proven most useful. Markets are naturally studied from the cooperative perspective since agents must act together—by matching or trading—to accomplish anything of value. The details of a market's operation are formidable and focusing on observed outcomes, rather than all contingent and counterfactual interactions, is a compelling and practical simplification.

Second, bracketing the within-period matching problem, a multiperiod economy involves sequential interactions with outcomes realized at multiple moments in time. On this front, traditionally non-cooperative approaches appear more successful. The logic of backward induction or subgame perfection is hard to ignore. Questions regarding dynamic consistency and (lack of) commitment over time have been most ably tackled with these tools in mind.

Our family of solutions, whose baseline is called perfect α-stability, threads together the above observations according to the following prototypical principle:

> A matching is α-stable in period t if (given the elapsed history) there is no coalition of agents who prefer an alternative period-t assignment given all plausible continuations of the market at the proposed alternative. A plausible continuation must be α-stable in period $t+1$. A perfect α-stable matching is α-stable in each period.

The above definition embeds three features. First, the solution is *cooperative* and focuses on outcomes that cannot be improved upon, or "blocked," by any coalition. Second, a *perfection requirement* and an intuitive recursive structure ensure the credibility of realized outcomes and potential blocking plans despite limited commitment. Third, because alternatives must be superior for "all plausible continuations," the definition incorporates a *robustness* notion similar to the α-core (Aumann and Peleg, 1960; Aumann, 1961). The α-core is well-suited for the study of economies with externalities (Scarf, 1971, p. 174). Although externalities are not formally in our model, they arise endogenously in a multiperiod economy, as explained below.

The definition stated above is incomplete. That is because the "plausibility" of a continuation remains undefined. We suggest three ways of defining this term that differ in—loosely speaking—agents' beliefs concerning market developments. Our proposals are hardly exhaustive and institutional or behavioral suppositions will motivate further variations on their themes. Perfect α-stability, our baseline, posits a cautious, "worst-case" disposition. A perfect α-stable matching exists, even when assignments are inter-temporal complements. The perfect α-core and perfect α^*-stability are refinements that increase agents' proclivity to block assignments by moderating beliefs. The perfect α-core constrains beliefs concerning the behavior of blocking coalition members; perfect α^*-stability rules out "incredible beliefs" reliant on dominated outcomes. Each solution allows for general coalitions to form and, surprisingly, effective blocking coalitions may involve only multiple agents from the same side of the market. Though novel in models of two-sided matching markets, actions by one-sided coalitions are empirically relevant. Strikes and cartel behavior are familiar examples.

Outline The next section reviews the literature and Section 3 introduces the model. Section 4 sketches the shortcomings of stability and the core, two standard static solutions. Sections 5, 6, and 7 introduce perfect α-stability, the perfect α-core, and perfect α^*-stability, respectively. An important feature of perfect α-stability is its adaptability to many related problems and applications. Section 8 extends our analysis to markets with arrivals and departures, transfers, many-to-one assignments, and an infinite time horizon. Appendices A and B collect all proofs and omitted examples.

2 Literature

The extension of Gale and Shapley's stability notion to a multiperiod or dynamic setting has drawn interest from many authors. Kadam and Kotowski (2018a,b) study a model closest to ours and propose a solution that is a weakening of Gale's (1978) sequential core (see below) or Becker and Chakrabarti's (1995) recursive core.[1] Though tractable, solutions in this class have three drawbacks. First, they are often empty. Second, they require blocking coalitions to disengage from the wider economy. Third, a blocking plan's credibility is often debatable. Damiano and Lam (2005) tackle the final shortcoming through their concepts of self-sustaining stability and strict self-sustaining stability. These solutions impose a coalition proofness (Bernheim et al., 1987) requirement on blocking plans. Damiano and Lam assume time-separable preferences, a restriction we do not impose.

A matching in our model only specifies realized allocations. Kurino (2020) and Doval (2018) define a "matching" as a complete contingent plan for the market. Liu (2018) studies a "matching process," which is similar.[2] This approach follows Corbae et al. (2003) who use a dynamic bilateral matching market to investigate questions in monetary economics. Corbae et al. (2003), Doval (2018), and Liu (2018) impose a perfection requirement in their solutions, a feature shared by our proposals. Kurino (2020) and Doval (2018) show that their solutions may be empty when applied to a multiperiod, one-to-one matching market. Their analyses are confined to different preference domains than we consider.[3]

The above studies mainly focus on one-to-one matching markets. Bando (2012) considers a multiperiod market where agents match with multiple counterparties each period. His solution adapts the usual stability definition from a static many-to-many matching market (Roth, 1984). An agent blocking a matching in period t takes his assignments in period $t' > t$ as given, an assumption not shared by our approach.

Pereyra (2013), Kennes et al. (2014, 2019), and Dur (2012) study a related class of assignment problems where agents' preferences and priorities to objects (typically seats at a school)

[1]Related to Gale (1978) is a literature on dynamic core concepts for incomplete market economies. Habis and Herings (2011) discuss several proposed definitions.

[2]Ali and Liu (2019) study "plans" and "conventions" in repeated games, which are also related.

[3]Kurino (2020) assumes time-separable preferences. Doval (2018) posits that agents match only one time and allows for stochastic arrivals. See also Thakral (2019). Altınok (2019) extends Doval's (2018) model to the case of many-to-one matching while Du and Livne (2016) examine a similar problem allowing for transfers.

vary with time and with prior assignments. They propose fairness criteria for their assignment problems, which can be interpreted as stability notions in a two-sided market. Their proposals involve institutional features of specific applications and differ from our analysis. Some studies examine matchings that form over multiple periods via an iterative process (Narita, 2018; Haeringer and Iehlé, 2019). Our model can be interpreted in this way, though we posit agents care about the sequence of (interim) assignments and not just the final-period match.[4] Diamantoudi et al. (2015) and Zhang and Zheng (2016) explore the role of commitment in multiperiod matching markets. Throughout we assume no commitment.

The solutions we study relate to the α-core (Aumann and Peleg, 1960; Aumann, 1961; Scarf, 1971).[5] They involve identifying outcomes that cannot be improved upon by any coalition, independently of others' contemporaneous actions. This logic's application to our problem is inspired by Sasaki and Toda's (1996) study of a one-period matching market with externalities. The conceptual parallel is the following. In Sasaki and Toda (1996), agents impose externalities on others when they match. In our setting, direct externalities are absent, but an agent's period-t assignment affects what outcomes can be stable in period $t' > t$. This endogenously introduces an externality. The complementary coalition's period-t assignments matter through this channel and motivate our adaptation of the α-core idea.[6]

Doval (2019) studies a dynamic matching market with stochastically arriving agents who form irrevocable assignments. Her solution concept, "dynamic stability," also relies upon pessimistic conjectures (Sasaki and Toda, 1996) to tackle the externality arising when agents delay matching.[7] Agents may match multiple times in our model. Thus, the nature of intertemporal externalities differs due to the complementarity or substitutability of successive assignments. Moreover, our solutions allow for general blocking coalitions while Doval (2019) focuses on pairwise blocking. This distinction is substantive, as explained in Section 6. A technical difference is also notable. Doval (2019) constructs dynamically stable matchings by

[4]For example, many school admission systems allow for appeals or reapplications. Interpreted within our model, the initial assignment is the period 1 matching. The final assignment, after appeals/reapplications are processed, is the period 2 matching. A participant will care about the sequence of assignments if, for example, there are transaction, opportunity, or psychological costs associated with the appeals/reapplications process.

[5]The α-core was proposed as a method for translating a strategic-form game into a cooperative game. This is not precisely the exercise pursued here since the within-period interaction in our model is a cooperative game.

[6]Corbae et al. (2003), Liu (2018), and Kurino (2020) acknowledge the dependence of a blocking coalition's payoff on others' actions through the change in history. They address this dependence by assuming non-blocking coalition members either become unmatched or continue with their prior assignment. This assumption implicitly imposes some structure on the market's within-period operation. By adapting the α-core's reasoning, our solution is robust to all within-period interactions among agents, which we do not model directly.

[7]"Dynamic stability" (Doval, 2019) weakens the solution examined by Doval (2018) by introducing pessimistic conjectures. Both solutions differ from "dynamic stability" as defined by Kadam and Kotowski (2018a,b).

examining a sequence of markets with successively longer time horizons. Our analysis relies on backward induction instead.

3 Model and Notation

Let $M=\{m_1,\dots,m_n\}$ and $W=\{w_1,\dots,w_{n'}\}$ be finite and disjoint sets of agents, labeled men and women, respectively. We denote generic agents by $i,j,k \in I := M\cup W$. Agents interact over $T<\infty$ periods. In every period, each man (woman) can be matched to one woman (man) or not matched to anyone. By convention, an unmatched agent is "matched to him/herself." Thus, the set of potential partners for $m\in M$ in period t is $W\cup\{m\}$. Symmetrically, the set of potential partners for $w\in W$ in period t is $M\cup\{w\}$.

A (multiperiod) *matching* $\mu\colon M\cup W \to (M\cup W)^T$ specifies each agent's partner in each period. It is a tuple $\mu=(\mu_1,\dots,\mu_T)$ of T (one-period) *assignments* that satisfy the following standard properties: (a) $\mu_t(m)\in W\cup\{m\}$ for all $m\in M$, (b) $\mu_t(w)\in M\cup\{w\}$ for all $w\in W$, and (c) $\mu_t(i)=j \implies \mu_t(j)=i$. The sets of assignments and matchings are $\mathscr{A}$ and $\mathscr{M}:=\mathscr{A}^T$, respectively.

Each agent has a strict preference over sequences of assigned partners. For each $m\in M$, $\succ_m$ is a strict preference defined over $(W\cup\{m\})^T$; $\succ_w$ is defined symmetrically for each $w\in W$. We write $x\succ_i y$ if agent i prefers the assignment sequence x over y and $x\succeq_i y$ if $x\succ_i y$ or $x=y$. Notwithstanding the lack of indifferences or externalities, this is the most general class of preferences applicable to our matching problem. Successive assignments may be complements or substitutes. Of course, history-independent preferences are a special case.

Notation We end this section by introducing notation used throughout the sequel. Given $K\subseteq I$ and $\mu_t\in\mathscr{A}$, let $\mu_t(K):=\bigcup_{i\in K}\mu_t(i)$. For any $K,K'\subseteq I$ and $\mathscr{A}'\subseteq\mathscr{A}$,

$$\mathscr{A}'(K,K') := \left\{\mu_t\in\mathscr{A}' \;\middle|\; \mu_t(K)=K'\right\}$$

is the set of assignments in $\mathscr{A}'$ where agents in K are assigned to agents in K'. When $K=K'$ and agents in coalition K are assigned only among themselves, we write $\mathscr{A}'(K):=\mathscr{A}'(K,K)$. The set of matchings where agents in coalition K match among themselves in each period is $\mathscr{M}(K):=\mathscr{A}(K)^T$. Given $\mathscr{A}'\subseteq\mathscr{A}$, $\mu_t\in\mathscr{A}'$, and $K\subseteq I$,

$$\mathscr{A}'(\mu_t|K) := \left\{\mu'_t\in\mathscr{A}' \;\middle|\; \mu'_t(i)=\mu_t(i)\ \forall i\in K\right\}$$

is the set of assignments in $\mathscr{A}'$ that coincide with μ_t on K, but possibly differ on $I \setminus K$.

The truncation of $\mu \in \mathscr{M}$ at period t is $\mu_{\leq t} \coloneqq (\mu_1, \ldots, \mu_t) \in \mathscr{M}_{\leq t}$. Its continuation from period t is $\mu_{\geq t} \coloneqq (\mu_t, \ldots, \mu_T) \in \mathscr{M}_{\geq t}$. We define $\mu_{<t} \in \mathscr{M}_{<t}$ and $\mu_{>t} \in \mathscr{M}_{>t}$ similarly. The matching $(\mu_{<t}, \mu'_t, \tilde{\mu}_{>t})$ means $(\mu_1, \ldots, \mu_{t-1}, \mu'_t, \tilde{\mu}_{t+1}, \ldots, \tilde{\mu}_T)$. The sets $\mathscr{M}_{<1}$ and $\mathscr{M}_{>T}$ are empty.

4 Stability and the Core

Stability and the core are the standard solutions applied to two-sided matching problems. The next definition applies irrespective of time horizon.

Definition 1. The coalition $K \subseteq I$ can *block* $\mu \in \mathscr{M}$ if there exists $\sigma \in \mathscr{M}(K)$ such that $\sigma(i) \succ_i \mu(i)$ for all $i \in K$.

A matching is *(pairwise) stable* if it cannot be blocked by either a single agent or a man-woman pair. A matching is in the *core* if it cannot be blocked by any nonempty coalition.

Stability and the core yield appealing predictions in a one-period economy. When $T = 1$, our model reduces to that of Gale and Shapley (1962); the stable set and the core coincide and are not empty. However, applying these solutions to a multiperiod market can be dissatisfying, especially if commitment is imperfect. (Under Definition 1, blocking happens ex ante.) A way to introduce imperfect commitment is to additionally allow blocking coalitions to form in any period conditional on the market's history.

Definition 2. The coalition $K \subseteq I$ can *block* $\mu \in \mathscr{M}$ *in period* t if there exists $\sigma_{\geq t} \in \mathscr{M}_{\geq t}(K)$ such that $(\mu_{<t}(i), \sigma_{\geq t}(i)) \succ_i \mu(i)$ for all $i \in K$.

A matching is in the *sequential core* if it cannot be blocked in any period by any nonempty coalition (Gale, 1978).

None of the above solutions offers entirely satisfactory predictions in a multiperiod matching market. Problems arise even if there is only one man and one woman.

Example 1. Consider a two-period market with one man and one woman.[8] Their preferences are:

$$\succ_m : wm, \underbrace{ww}_{\sigma}, \underbrace{\mathbf{mm}}_{\mu}, mw \qquad \succ_w : \underbrace{mm}_{\sigma}, \underbrace{\mathbf{ww}}_{\mu}, mw, wm.$$

[8]This example is adapted from Kadam and Kotowski (2018a). Hatfield and Kominers (2017) propose a related example of a doctor and a hospital contracting morning and afternoon shifts.

We state preferences by listing assignment sequences in preferred order. We omit the usual brackets and commas for clarity. Matchings are identified by highlighting the relevant assignments. In this case, both agents prefer to match for two periods (σ) rather than remain unmatched (μ). However, m's most-preferred option is to match with w only for period 1.

Some reflection suggests that μ, where the agents are unmatched, is this example's likely outcome. Both agents would rather match for two periods, but m prefers to be unmatched in period 2 conditional on matching with w in period 1 ($w\,m \succ_m w\,w$). If m refuses to continue the matching with w after period 1, w is worse off than had she not matched with m at all ($w\,w \succ_w m\,w$). Presumably, she would anticipate this possibility and shun σ entirely.

Neither the core nor the sequential core select μ in this example. The matching σ is this economy's unique pairwise stable and core matching. Its flaws were explained above. The sequential core is empty since m can block σ in period 2.

5 Perfect α-Stability

In this section we introduce perfect α-stability, our baseline solution. We start by formalizing the principle stated in the introduction to define cautious α-blocking. The definition is recursive and we provide a step-by-step explanation of its components and an example below.

Definition 3. The coalition $K \subseteq I$ can *cautiously α-block $\mu \in \mathcal{M}$ in period t* if there exists $\sigma_t \in \mathcal{A}(K)$ such that $\sigma_t(i) \neq \mu_t(i)$ for some $i \in K$ and for each $i \in K$,

$$\big(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)\big) \succ_i \mu(i) \quad \text{for all} \quad \sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathcal{A}(\sigma_t|i)} \mathcal{S}\big((\mu_{<t}, \tilde{\sigma}_t)\big) \tag{1}$$

where $\mathcal{S}((\mu_{<t}, \tilde{\sigma}_t))$ is the set of continuations $\sigma_{>t} \in \mathcal{M}_{>t}$ such that $(\mu_{<t}, \tilde{\sigma}_t, \sigma_{>t})$ cannot be cautiously α-blocked in any period $t' > t$.

A matching is *perfect α-stable* if it cannot be cautiously α-blocked in any period by any nonempty coalition.

To unpack Definition 3, it is best to start in period T. Coalition K can cautiously α-block μ in period T if and only if there exists $\sigma_T \in \mathcal{A}(K)$ such that $(\mu_{<T}(i), \sigma_T(i)) \succ_i (\mu_{<T}(i), \mu_T(i))$ for all $i \in K$. The next lemma is an immediate implication. Its proof is omitted.

Lemma 1. *If $\mu = (\mu_1, \ldots, \mu_T)$ cannot be cautiously α-blocked in period T, then μ_T is a core assignment in a one-period economy where the strict preference P_i of each agent i satisfies*

$$j P_i k \iff (\mu_{<T}(i), j) \succ_i (\mu_{<T}(i), k). \tag{2}$$

Lemma 1 implies that perfect α-stability reduces to (pairwise) stability in a one-period market.

Cautious α-blocking in period $t < T$ is more intricate. An assignment $\sigma_t \in \mathscr{A}(K)$ must be preferred by each member of coalition K given all plausible market continuations, the latter half of expression (1). Per Definition 3, agent i considers the continuation $\sigma_{>t}$ plausible if it cannot be cautiously α-blocked in the future given the market's history ($\mu_{<t}$) and given some period-t assignment ($\widetilde{\sigma}_t$) that preserves agent i's assignment at σ_t, i.e., $\widetilde{\sigma}_t \in \mathscr{A}(\sigma_t|i) = \{\sigma'_t \in \mathscr{A} | \sigma'_t(i) = \sigma_t(i)\}$. A cautiously α-blocked matching is unlikely to arise since a coalition has a robust incentive to depart from its prescriptions. A blocking agent benefits whatever else happens in period t.

Example 1 (Continued)**.** Recall the market with one man and one woman. The matching σ is not perfect α-stable. It can be blocked by m in period 2. The matching μ is perfect α-stable. It cannot be blocked in period 2 since both agents prefer to remain unmatched if unmatched in period 1. Likewise, μ cannot be cautiously α-blocked in period 1. To block μ in period 1, m and w must pursue $\sigma_1 \in \mathscr{A}(\{m, w\})$ where $\sigma_1(m) = w$ and $\sigma_1(w) = m$. For each $i \in \{m, w\}$, $\mathscr{A}(\sigma_1|i) = \{\sigma_1\}$. Conditional on σ_1, there is one assignment (i.e., continuation) that cannot be blocked in period 2: $\mathscr{S}(\sigma_1) = \{\sigma'_2\}$ where $\sigma'_2(m) = m$ and $\sigma'_2(w) = w$. However, $\mu(w) = (w, w) \succ_w (m, w) = (\sigma_1(w), \sigma'_2(w))$. Thus, w is unwilling to block in period 1.

Several features distinguish perfect α-stability from the solutions surveyed in Section 2. First, a blocking coalition does not disengage from the wider economy forever. Second, the credibility of a blocking action initiated in period t is assured as agents anticipate its continuation (in periods $t' > t$) to be stable. Third, members of a blocking coalition do not take their original future assignments as given. Instead, they anticipate stable assignments given their new arrangement. Fourth, the existence of a perfect α-stable outcome does not depend on a narrow preference domain. This feature is important as inter-temporal complementarities, switching costs, and other history-dependencies arise in many applications.

Theorem 1. *The set of perfect α-stable matchings is not empty.*

Despite yielding a nonempty solution, perfect α-stability is not a "weak" concept per se. Example 1 demonstrates its logical independence from the core, which may in general be empty.[9] Examples B.1–B.3 in Appendix B show that it is not weaker than the sequential core or the stability concepts of Damiano and Lam (2005), Kennes et al. (2014), and Kadam and Kotowski (2018a). It is also not weaker than Kurino's (2020) proposal, which is weaker than pairwise stability in a one-period economy.

[9]Kadam and Kotowski (2018a) provide an example of a two-period economy with an empty core.

The next two corollaries follow from the proof of Theorem 1. First, Definition 3 does not restrict a blocking coalition's composition. Nevertheless, cautious α-blocking reduces to blocking by singletons or pairs.

Corollary 1. *A coalition can cautiously α-block $\mu \in \mathscr{M}$ in period t if and only if μ can be cautiously α-blocked in period t by either a single agent or a man-woman pair.*

Second, perfect α-stability captures an important dimension of subgame perfection.

Corollary 2. *If $\mu = (\mu_1, \ldots, \mu_T)$ is perfect α-stable, then $\mu_{\geq t} = (\mu_t, \ldots, \mu_T)$ is perfect α-stable in the "submarket" beginning in period t where the preference P_i of each agent i satisfies $(x_t, \ldots, x_T) P_i (y_t, \ldots, y_T) \iff (\mu_{<t}(i), x_t, \ldots, x_T) \succ_i (\mu_{<t}(i), y_t, \ldots, y_T)$.*

Unlike a subgame perfect strategy profile in a non-cooperative game, a perfect α-stable matching does not define particular outcomes in counterfactual histories. This distinguishes it from solutions studied by Corbae et al. (2003), Doval (2018), Kurino (2020), and others.

The importance of timing and incentives in a sequential market means that many properties of stable matchings in a one-period economy no longer apply when $T \geq 2$. For example, a perfect α-stable matching may not be Pareto optimal (Example 1).[10] A one-period economy also has an "optimal stable matching" for each side of the market (Gale and Shapley, 1962). That is, among all stable matchings, all men (women) agree which is best. This alignment of interest does not extend to the perfect α-stable set (Example B.4 in Appendix B).[11]

6 The Perfect α-Core

When coalition K cautiously α-blocks μ in period t with σ_t, agent $i \in K$ considers all $\widetilde{\sigma}_t \in \mathscr{A}(\sigma_t | i)$ as possible period-t outcomes. This set includes assignments where members of the blocking coalition other than i and his partner are not matched according to σ_t. Considering these possibilities can be sometimes prudent. Agent i's information may be limited due to imperfect monitoring or he may distrust others. Nevertheless, removing this layer of cautious, seemingly non-cooperative, beliefs from Definition 3 is tempting. Doing so involves replacing $\mathscr{A}(\sigma_t | i)$ in (1) with $\mathscr{A}(\sigma_t | K)$. Now, agent $i \in K$ is confident that the assignment of each $j \in K$ is fixed at $\sigma_t(j)$. This weakening of Definition 3 results in a stronger solution.

[10] The possible inefficiency of stable or equilibrium outcomes in a dynamic model is not unusual. See Damiano and Lam (2005) or Kurino (2020) for other examples. In Sasaki and Toda (1996), a Pareto optimal stable matching always exists. Thus, our results are not implied by their analysis despite the connection noted in Section 2.

[11] The absence of an optimal stable matching implies that the perfect α-stable set is not a lattice under the common preference ordering. See Roth and Sotomayor (1990) for this property's definition and implications.

Definition 4. The coalition $K \subseteq I$ can *α-block $\mu \in \mathscr{M}$ in period t* if there exists $\sigma_t \in \mathscr{A}(K)$ such that $\sigma_t(i) \neq \mu_t(i)$ for some $i \in K$ and for each $i \in K$,

$$\left(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)\right) \succ_i \mu(i) \quad \text{for all} \quad \sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|K)} \mathscr{C}\left((\mu_{<t}, \tilde{\sigma}_t)\right) \tag{3}$$

where $\mathscr{C}((\mu_{<t}, \tilde{\sigma}_t))$ is the set of continuations $\sigma_{>t} \in \mathscr{M}_{>t}$ such that $(\mu_{<t}, \tilde{\sigma}_t, \sigma_{>t})$ cannot be α-blocked in any period $t' > t$.

The *perfect α-core* is the set of matchings that cannot be α-blocked in any period by any nonempty coalition.

Theorem 2. *Every perfect α-core matching is perfect α-stable.*

If μ cannot be α-blocked in period T, then μ_T is a core matching in a one-period economy where (2) defines agents' preferences. Thus, the perfect α-core equals the core and the perfect α-stable set in a one-period economy. In general, however, the perfect α-core can be a strict subset of the perfect α-stable set. Moreover, α-blocking by a coalition does not reduce to α-blocking by singletons or man-woman pairs (cf. Corollary 1). In the following example, a matching is blocked by a coalition of two women. Neither woman can α-block the outcome independently. Though superfluous in typical (static) matching market models, one-sided collective actions are economically relevant, with strikes being a familiar example.[12] Definitions 3 and 4 do not restrict a blocking coalition's membership. Thus, the relevance of non-singleton and non-pairwise coalitions is driven entirely by agents' beliefs concerning others' matches at a blocking assignment σ_t, the sets $\mathscr{A}(\sigma_t|i)$ and $\mathscr{A}(\sigma_t|K)$ in (1) and (3).

Example 2. Consider a two-period economy with two men and two women. Their preferences are:

$$\succ_{m_1}: \underbrace{w_2 w_2}_{\sigma}, \underbrace{\mathbf{m_1 w_1}}_{\mu}, m_1 m_1, \ldots \qquad \succ_{w_1}: \underbrace{\mathbf{w_1 m_1}}_{\mu}, \underbrace{m_2 m_2}_{\sigma}, w_1 w_1, \ldots$$

$$\succ_{m_2}: \underbrace{w_1 w_1}_{\sigma}, \underbrace{\mathbf{m_2 w_2}}_{\mu}, m_2 m_2, \ldots \qquad \succ_{w_2}: \underbrace{\mathbf{w_2 m_2}}_{\mu}, \underbrace{m_1 m_1}_{\sigma}, w_2 w_2, \ldots$$

There are two perfect α-stable matchings, μ and σ. Only μ is in the perfect α-core. Neither woman can α-block σ alone. Instead, σ is α-blocked by a coalition of *both* women being unmatched in period 1.[13]

[12] In the case of a matching between men and women, Aristophanes' *Lysistrata* provides another example.

[13] An analogue of the blocking action in this example has been documented among drivers using ride-sharing platforms. Drivers collectively refuse rides to improve terms for subsequent matches. Quoting a news report:

> Every night, several times a night, Uber and Lyft drivers at Reagan National Airport simultaneously

The perfect α-core may be empty (Example B.5 in Appendix B). Nonemptiness is assured on some preference domains and we examine two cases of interest.

History Independence When preferences are history independent, an agent's assessment of a continuation at period t does not depend on prior assignments. Formally, $\succ_i$ is *history independent* if for each t and all $x = (x_1, \ldots, x_T)$ and $y = (y_1, \ldots, y_T)$, $(x_{<t}, x_{\geq t}) \succ_i (x_{<t}, y_{\geq t}) \implies (y_{<t}, x_{\geq t}) \succ_i (y_{<t}, y_{\geq t})$. For example, the preference

$$jk \succ_i kk \succ_i kj \succ_i jj \tag{4}$$

is history independent. Agent k is the superior period-2 assignment irrespective of the period-1 match. Time-separable preferences, often defined with an additively-separable utility function, are history independent.[14] As noted in Section 2, history independence and time-separability are common assumptions in the literature.

Theorem 3. *If each agent's preference is history independent, then the perfect α-core is not empty and equals the perfect α-stable set.*

Sequential Sacrifice Aversion A limitation of history independent preferences is that they cannot capture inter-temporal complementarities, switching costs, or status quo bias, all common phenomena (Samuelson and Zeckhauser, 1988). An example of a preference exhibiting these features in our model is

$$jj \succ_i jk \succ_i kk \succ_i kj. \tag{5}$$

In (5), agent i prefers to continue a matching with j conditional on matching with j in period 1. However, conditional on matching with k in period 1, continuing that assignment is best. Clearly, (5) is not history independent.

A class of preferences that includes (5) and allows for inter-temporal complementarities is

turn off their ride share apps for a minute or two to trick the app into thinking there are no drivers available—creating a price surge. When the fare goes high enough, the drivers turn their apps back on and lock into the higher fare. (Sweeney, 2019)

[14] Let $\mathcal{T} = \{1, \ldots, T\}$ be the set of period indices. Agent i's preference $\succ_i$ is *time-separable* if for each subset of period indices, $\mathcal{T}' \subseteq \mathcal{T}$, and all $x = (x_{\mathcal{T}'}, x_{\mathcal{T} \setminus \mathcal{T}'})$ and $y = (y_{\mathcal{T}'}, y_{\mathcal{T} \setminus \mathcal{T}'})$, $(x_{\mathcal{T}'}, x_{\mathcal{T} \setminus \mathcal{T}'}) \succ_i (x_{\mathcal{T}'}, y_{\mathcal{T} \setminus \mathcal{T}'}) \implies (y_{\mathcal{T}'}, x_{\mathcal{T} \setminus \mathcal{T}'}) \succ_i (y_{\mathcal{T}'}, y_{\mathcal{T} \setminus \mathcal{T}'})$. For example, $jk \succ_i kk \succ_i jj \succ_i kj$ is time-separable. Time-separability implies history independence. The converse is not true. The preference in (4) is not time-separable.

the following. The preference $\succ_i$ is *sequentially sacrifice averse (SSA)*[15]

$$(x_{<t}, x_t, \dots, x_T) \succeq_i (x_{<t}, j, \dots, j) \implies (x_{<t}, x_{t'}, \dots, x_{t'}) \succeq_i (x_{<t}, j, \dots, j)$$

for every $t' = t, \dots, T$. In words, if a sequentially sacrifice averse agent prefers a sequence of assignments over a constant assignment (starting in period t), then the sequence cannot involve any assignment that is a worse long-term outcome than the constant matching. Thus, the agent is averse to even a moment of sacrifice.

Theorem 4. *If each agent's preference satisfies SSA, then the perfect α-core is not empty.*

7 Perfect α^*-Stability

A potential drawback of perfect α-stability is that a matching may be rationalized by "incredible beliefs." In this section, we propose a refinement to address this peculiarity, which is illustrated by the following example. To limit repetition, the refinement's analogue for the perfect α-core is only sketched in Remark 1 below.

Example 3. Consider a two-period economy with two men and two women. Their preferences are:

$$\succ_{m_1}: \underbrace{\mathbf{w_1 w_1}}_{\mu,\sigma}, m_1 m_1, \dots \qquad \succ_{w_1}: \underbrace{\mathbf{m_1 m_1}}_{\mu,\sigma}, w_1 w_1, m_2 m_2, \dots$$

$$\succ_{m_2}: \underbrace{w_2 w_2}_{\sigma}, \underbrace{\mathbf{m_2 w_2}}_{\mu}, m_2 m_2, w_1 w_1, \dots \qquad \succ_{w_2}: \underbrace{\mathbf{w_2 m_2}}_{\mu}, \underbrace{m_2 m_2}_{\sigma}, w_2 w_2, \dots$$

The matchings μ and σ constitute this market's perfect α-stable set. Agent w_2 may wish to (cautiously) α-block σ by remaining unassigned in period 1. If w_2 is unassigned in period 1, w_1 and m_2 could match in period 1 instead. But then, m_2 and w_1 will prefer to match again in period 2. Thus, w_2 would be unmatched in both periods, an outcome that is worse than $\sigma(w_2) = m_2 m_2$. Hence, w_2 is content at σ and the matching is perfect α-stable.

Two facts suggest w_2's hesitation about blocking σ should be unfounded. First, the above logic requires w_1 and m_1 to forgo their most-preferred outcome (as if to "punish" w_2). And second, it also assumes w_1 and m_2 put aside their individual rationality. For instance, the *best* that w_1 can achieve after matching with m_2 in period 1 is $m_2 m_2$, which is *worse* than had she

[15] Kadam and Kotowski (2018b) introduce a weaker preference restriction that they call sacrifice aversion. The definitions are equivalent when $T \leq 2$.

not matched with him at all. Together, these reasons suggest that w_2 should ignore the contingency where w_1 and m_2 match in period 1 when evaluating the implications of remaining unmatched in period 1. Thus, we might regard σ with some skepticism.

Motivated by Example 3, a natural refinement of perfect α-stability posits that agents ignore "dominated" outcomes when evaluating a blocking assignment. Eliminating dominated strategies from rationalizing beliefs has been considered by Bernheim (1984), Pearce (1984), Ambrus (2006), and Liu et al. (2014), among others. We use this idea to thin the set $\mathscr{A}(\sigma_t|i)$ in Definition 3 by removing dominated assignments for agents $j \neq i$. We call the resulting blocking notion cautious α^*-blocking.

Following Definition 3, cautious α^*-blocking is defined recursively, a fact incorporated into the underlying dominance notion. Accordingly, let $\mathscr{S}^*(\mu_{\leq t})$ be the set of all $\mu'_{>t} \in \mathscr{M}_{>t}$ such that the matching $(\mu_{\leq t}, \mu'_{>t})$ cannot be cautiously α^*-blocked in any period $t' > t$ by any nonempty coalition.

Definition 5. Let $\mu_{<t} \in \mathscr{M}_{<t}$, $\sigma_t \in \mathscr{A}$, and $K \subseteq I$. Assignment $\nu_t \in \mathscr{A}(\sigma_t|K)$ is *considered by K to be cautiously α-dominated at* $(\mu_{<t}, \sigma_t)$ if there exists a nonempty $J \subseteq I \setminus K$ and $\widehat{\nu}_t \in \mathscr{A}(\nu_t|K) \cap \mathscr{A}(J)$ such that for all $j \in J$, all

$$\widehat{\nu}_{>t} \in \bigcup_{\widehat{\nu}'_t \in \mathscr{A}(\widehat{\nu}_t|j)} \mathscr{S}^*\big((\mu_{<t}, \widehat{\nu}'_t)\big), \tag{6a}$$

and all

$$\nu_{>t} \in \bigcup_{\nu'_t \in \mathscr{A}(\nu_t|j)} \mathscr{S}^*\big((\mu_{<t}, \nu'_t)\big), \tag{6b}$$

$(\mu_{<t}(j), \widehat{\nu}_t(j), \widehat{\nu}_{>t}(j)) \succ_j (\mu_{<t}(j), \nu_t(j), \nu_{>t}(j))$.

Intuitively, ν_t is dominated if a group of agents J has an assignment that is superior to ν_t in every contingency (accounting for continuation matchings themselves being stable). Let $\mathscr{A}^*_{\mu_{<t}}(\sigma_t|K) \subseteq \mathscr{A}(\sigma_t|K)$ be the set of assignments that are *not* considered by K to be cautiously α-dominated at $(\mu_{<t}, \sigma_t)$. This set is not empty (Lemma A.2 in Appendix A). The next definition mirrors Definition 3 with $\mathscr{A}^*_{\mu_{<t}}$ and $\mathscr{S}^*$ replacing $\mathscr{A}$ and $\mathscr{S}$, respectively.

Definition 6. The coalition K can *cautiously α^*-block $\mu \in \mathscr{M}$ in period t* if there exists $\sigma_t \in \mathscr{A}(K)$ such that $\sigma_t(i) \neq \mu_t(i)$ for some $i \in K$ and for each $i \in K$,

$$\big(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)\big) \succ_i \mu(i) \quad \text{for all} \quad \sigma_{>t} \in \bigcup_{\widetilde{\sigma}_t \in \mathscr{A}^*_{\mu_{<t}}(\sigma_t|i)} \mathscr{S}^*\big((\mu_{<t}, \widetilde{\sigma}_t)\big)$$

where $\mathscr{S}^*((\mu_{<t},\widetilde{\sigma}_t))$ is the set of continuations $\sigma_{>t} \in \mathscr{M}_{>t}$ such that $(\mu_{<t},\widetilde{\sigma}_t,\sigma_{>t})$ cannot be cautiously α^*-blocked in any period $t' > t$.

A matching is *perfect α^*-stable* if it cannot be cautiously α^*-blocked in any period by any nonempty coalition.

Theorem 5. *The set of perfect α^*-stable matchings is not empty.*

Theorem 6. *Every perfect α^*-stable matching is perfect α-stable.*

Though equivalent in a one-period economy, perfect α^*-stability is stronger than perfect α-stability. In Example 3, σ is not perfect α^*-stable.

Remark 1. An analogous refinement of the perfect α-core, the perfect α^*-core, can be proposed. Call the associated blocking notion α^*-blocking. In Definition 5, replace $\mathscr{S}^*$ in (6a) and (6b) with the set of continuations that cannot be α^*-blocked in any future period. Also, replace $\mathscr{A}(\widehat{\nu}_t|j)$ in (6a) and (6b) with $\mathscr{A}(\widehat{\nu}_t|J)$. This ensures that the premise of α-blocking (the assignment of *all* blocking coalition members is fixed) is reflected in the dominance criterion. The definitions of α^*-blocking and the perfect α^*-core follow accordingly.

Remark 2. Perfect α^*-stability is but one possible refinement of perfect α-stability. Any alternative only requires a sensible trimming of $\mathscr{A}(\sigma_t|i)$ along with an accompanying adjustment to $\mathscr{S}(\cdot)$. Institutional features may help formulate refinements in applications. In his extension of Sasaki and Toda's (1996) analysis of a matching market with externalities, Hafalir (2008) proposes a different refinement of that model's analogue of $\mathscr{A}(\sigma_t|i)$. Hafalir's proposal adds assignments to a set of possibilities. Our proposal shrinks a large initial set.

8 Extensions

The framework afforded by perfect α-stability can accommodate many extensions.

Arrivals and Departures The set of market participants often changes with time. Some agents might be long-lived; others may be active for only one period. The model above requires two changes to admit arrivals and departures. First, replace the set of assignments $\mathscr{A}$ with a date-specific set of admissible assignments $\mathscr{A}_t$. Each $\mu_t \in \mathscr{A}_t$ is an assignment among agents present in period t, say $I_t = M_t \cup W_t$. The set $\mathscr{A}_t$ passes through to the rest of the model—the set of matchings becomes $\mathscr{M} = \times_{t=1}^{T}\mathscr{A}_t$, the set $\mathscr{A}(K)$ becomes $\mathscr{A}_t(K)$, and so on. Second, restrict agents' preferences to the periods when they are in the market.[16] The rest of the analysis is unchanged.

Historical Dependencies Above we assumed that an agent's preferences may depend on his prior assignments. In practice, the prior assignments of *others* matter too. For example, many employers value workers with experience in the same industry and colleges care about applicants' prior education. It is simple to verify that a perfect α-stable matching exists when agents' preferences over continuations from period t (the set $\mathscr{M}_{\geq t}$) depend on the market-wide history of assignments up to period t, $\mu_{<t}(\cdot)$. Only a minor amendment to the proof of Theorem 1—additionally conditioning on $\mu_{<t}(\cdot)$—is required.

Transfers Perfect α-stability provides a template that can be adapted to many problems. For example, consider a multiperiod version of Shapley and Shubik's (1971) matching market with transfers. The market has two sides and only bilateral matches generate surplus. Following convention, we now say that an *assignment for period t* is a matrix $x_t = [x_t^{mw}]$ such that (a) $\sum_{w\in W} x_t^{mw} \leq 1$ for all $m \in M$, (b) $\sum_{m\in M} x_t^{mw} \leq 1$ for all $w \in W$, and (c) $x_t^{mw} \in \{0,1\}$ for all m and w. If $x_t^{mw} = 1$, then m and w are matched and generate $v_t^{mw} \geq 0$ in surplus. Let $\mathscr{X}$ be the set of assignment matrices and let $\mathscr{X}(K)$ be the set of assignment matrices where each agent in coalition K is matched with another member of K or is unassigned.

An *imputation for period t* is a division of the period-t surplus. It is a vector $u_t = (u_t^i)_{i\in I} \in \mathbb{R}^{|I|}$ of utility values, one for each agent. An imputation u_t is *feasible for coalition K at $x_t \in \mathscr{X}(K)$* if $\sum_{i\in K} u_t^i \leq \sum_{(m,w)\in M\cap K\times W\cap K} x_t^{mw} v_t^{mw}$.[17] An imputation is *feasible* if it is feasible for the grand coalition $K = I$. A *period-t outcome* $\gamma_t = (x_t, u_t)$ consists of an assignment and a feasible imputation. An *outcome* $\gamma = (\gamma_1, \ldots, \gamma_T)$ is a tuple of T one-period outcomes.

Each agent's preference is defined over sequences of per-period utilities. An important special case assumes that agent i's preference is represented by the function

$$U_i\left(u_1^i, \ldots, u_T^i\right) = \sum_{t=1}^{T} \delta^{t-1} u_t^i. \tag{7}$$

where $\delta \in (0,1)$ is a discount factor.

We adapt prior ideas in the natural way to this model. For each $\gamma_t = (x_t, u_t)$, let $\mathscr{A}(\gamma_t|K)$ be the set of all period-t outcomes $\tilde{\gamma}_t = (\tilde{x}_t, \tilde{u}_t)$ such that $\tilde{u}_t$ is feasible at $\tilde{x}_t$ and the period-t

[16] For example, an agent active from period t to period $t+2$ would have preferences defined over triples of the form ijk, his period t, $t+1$, and $t+2$ assignments.

[17] By convention $\sum_{i\in J}(\cdot) = 0$ if $J = \varnothing$.

outcome of members of K equals γ_t, i.e., $\tilde{x}_t^{mw} = x_t^{mw}$ for all $m, w \in K$ and $\tilde{u}_t^i = u_t^i$ for all $i \in K$.

Definition 7. The coalition $K \subseteq I$ can *cautiously α-block the outcome γ in period t* if there exists $\sigma_t = (y_t, s_t) \in \mathscr{X}(K) \times \mathbb{R}^{|I|}$ such that (a) s_t is feasible for coalition K at y_t, (b) the assignment or imputation of some $i \in K$ changes, and (c) for each $i \in K$,

$$\left(u^i_{<t}, s^i_t, s^i_{>t}\right) \succ_i \left(u^i_1, \ldots, u^i_T\right) \quad \text{for all} \quad \sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|i)} \mathscr{S}\left((\gamma_{\le t}, \tilde{\sigma}_t)\right),$$

where s^i_t and $s^i_{>t}$ are the utilities for agent i at σ_t and $\sigma_{>t}$, respectively, and $\mathscr{S}((\gamma_{<t}, \tilde{\sigma}_t))$ is the set of continuations $\sigma_{>t}$ such that $(\gamma_{<t}, \tilde{\sigma}_t, \sigma_{>t})$ cannot be cautiously α-blocked in any period $t' > t$ by any nonempty coalition.

An outcome is *perfect α-stable* if it cannot be cautiously α-blocked in any period by any nonempty coalition. The definition of (non-cautious) α-blocking mimics Definition 7 with $\mathscr{A}(\cdot|K)$ and $\mathscr{C}(\cdot)$ replacing $\mathscr{A}(\cdot|i)$ and $\mathscr{S}(\cdot)$, respectively. The *perfect α-core* is the set of outcomes that cannot be α-blocked in any period by any nonempty coalition.

Theorem 7. *Suppose each agent's preference is represented by* (7)*. The set of perfect α-stable outcomes is not empty and equals the perfect α-core.*

Many-to-One Assignments The original many-to-one assignment model concerns college admissions (Gale and Shapley, 1962). One interpretation of this model in a multiperiod setting is that students pursue "education plans" by attending a sequence of schools.[18] For example, a student may plan to enroll in community college and then transfer to a university. About 26,500 California community college transfer students were admitted to the University of California system in 2019 (UCOP, 2019) and the system has particular programs, such as the UC Transfer Admission Guarantee and Transfer Admissions Pathways, to facilitate transfers.

Adopting standard nomenclature, let $S = \{s_1, \ldots, s_n\}$ and $C = \{c_1, \ldots, c_{n'}\}$ be the sets of students and colleges, respectively. College c has capacity q^c_t in period t. Now, a feasible period-t assignment is a function $\mu_t \colon C \cup S \to C \cup 2^S$ such that (a) $\mu_t(s) \in C \cup \{s\}$ for all $s \in S$, (b) $\mu_t(c) \in 2^S$ and $|\mu_t(c)| \le q^c_t$ for all $c \in C$, and (c) for $s \in S$ and $c \in C$, $\mu_t(s) = c \iff s \in \mu_t(c)$. Let $\mathscr{A}_t$ be the set of feasible period-t assignments. A matching $\mu = (\mu_1, \ldots, \mu_T) \in \mathscr{M} = \times_{t=1}^T \mathscr{A}_t$ is a tuple of T feasible assignments. Each student's preference $\succ_s$ is defined over $(C \cup \{s\})^T$. Each college's preference $\succ_c$ is defined over sequences of enrolled classes (sets of students).

[18]Each period can be an academic year. Enrolled students are typically assured "admission" to the same school for the following year. This practice can be encoded in the schools' preferences.

We can extend standard solutions to the static college admissions problem to this dynamic setting via the template provided by perfect α-stability. For brevity, we define only the analogue of group stability (Roth and Sotomayor, 1990, Definition 5.4).[19]

Definition 8. Coalition $K \subseteq C \cup S$ can *cautiously* α*-block* $\mu \in \mathscr{M}$ in period t if there exists $\sigma_t \in \mathscr{A}_t$, $\sigma_t(i) \neq \mu_t(i)$ for some $i \in K$, such that

(a) $\sigma_t(s) \in (K \cap C) \cup \{s\}$ for all $s \in K \cap S$,

(b) $\sigma_t(c) \subseteq (K \cap S) \cup \mu_t(c)$ for all $c \in K \cap C$, and

(c) for each $i \in K$, $(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)) \succ_i \mu(i)$ for all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}_t(\sigma_t|i)} \mathscr{S}((\mu_{\leq t}, \tilde{\sigma}_t))$ where $\mathscr{S}((\mu_{\leq t}, \tilde{\sigma}_t))$ is the set of continuations $\sigma_{>t} \in \mathscr{M}_{>t}$ such that $(\mu_{<t}, \tilde{\sigma}_t, \sigma_{>t})$ cannot be cautiously α-blocked by any nonempty coalition in any period $t' > t$.

Point (a) says each blocking student matches with a college in the blocking coalition or is unassigned. Point (b) says each blocking college may enroll new students while possibly keeping or dropping some prior assignments. (The college's capacity constraint is satisfied since $\sigma_t \in \mathscr{A}_t$.) Point (c) is the usual decision criterion for cautious α-blocking.

A matching is *perfect* α*-group stable* if it cannot be cautiously α-blocked by any coalition in any period. In a one-period economy, this definition reduces to group stability and a restriction on the colleges' preferences is needed to ensure that a group stable outcome exists. A common assumption is responsiveness (Roth, 1985). A ranking P_c of 2^S is *responsive* if for each $x_t, y_t \in 2^S$ such that $i \notin y_t$, $j \in y_t$, and $x_t = y_t \cup \{i\} \setminus \{j\}$, $x_t P_c y_t \iff \{i\} P_c \{j\}$. Call the preference of college c *conditionally responsive* if for all t and $\mu_{<t} \in \mathscr{M}_{<t}$ there exists a strict and responsive ranking of 2^S, $P_c^{\mu_{<t}}$, such that $(\mu_{<t}(c), x_t, x_{>t}) \succ_c (\mu_{<t}(c), y_t, y_{>t}) \iff x_t P_c^{\mu_{<t}} y_t$.

Theorem 8. *Suppose colleges have conditionally responsive preferences. (No additional restrictions are placed on the students' preferences.)*

(a) If coalition K can cautiously α-block $\mu \in \mathscr{M}$ in period t, then μ can be cautiously α-blocked in period t by either a student alone, a college alone, or a college-student pair.

(b) The set of perfect α-group stable matchings is not empty.

[19] Roth and Sotomayor (1990) describe several solution concepts for the college admissions model—stability, group stability, and the core with weak or strict domination. Stability limits blocking coalitions to singletons or college-student pairs. It is equivalent to group stability when colleges have responsive preferences.

Infinite Time Horizon As a final extension, we revisit our original model but allow for an infinite time horizon. Since there is no final period on which to anchor beliefs about future outcomes (required for backward induction), we generalize our solution by explicitly defining such beliefs. A consistency requirement closes the definition.

Definition 9. The matching μ is *perfect* $\overline{\alpha}$*-stable* if there exists a mapping $\overline{\mathscr{S}}: \bigcup_{t=0}^{\infty} \mathscr{M}_{\le t} \to \bigcup_{t=0}^{\infty} \mathscr{M}_{>t}$ such that:

(a) No coalition $K \subseteq I$ can *cautiously* $\overline{\alpha}$*-block* μ (given $\overline{\mathscr{S}}(\cdot)$) in any period t. That is, there does not exist a coalition $K \subseteq I$, period t, and $\sigma_t \in \mathscr{A}(K)$ such that $\sigma_t(i) \neq \mu_t(i)$ for some $i \in K$ and for each $i \in K$, $(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)) \succ_i \mu(i)$ for all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|i)} \overline{\mathscr{S}}((\mu_{<t}, \tilde{\sigma}_t))$.

(b) $\overline{\mathscr{S}}(\cdot)$ is *consistent.* That is, $\sigma_{>t} \in \overline{\mathscr{S}}(\sigma_{\le t})$ if and only if $(\sigma_{\le t}, \sigma_{>t})$ cannot be cautiously $\overline{\alpha}$-blocked (given $\overline{\mathscr{S}}(\cdot)$) in any period $t' > t$ by any coalition.

In the preceding definition, the function $\overline{\mathscr{S}}(\cdot)$ represents agents' common beliefs about plausible market continuations. In a finite-horizon market, $\overline{\mathscr{S}}(\cdot)$ is pinned down by backward induction and perfect $\overline{\alpha}$-stability reduces to perfect α-stability.

Payoff restrictions are necessary to ensure the existence of an equilibrium in infinite horizon games. The situation here is no different. We say that agent i *eventually prefers constant assignments* if there exists a T_i such that for all $t \ge T_i$, $(x_{<t}, x_{\ge t}) \succeq_i (x_{<t}, i, i, \ldots) \implies x_{\ge t} = (j, j, \ldots)$. Every (finite) T-period market satisfies this condition (i.e., $T_i = T$ for each i). Intuitively, the condition says that agents do not value changes in assignments that occur sufficiently far in the future. The condition allows us to define $\overline{\mathscr{S}}(\mu_{\le t})$ consistently for sufficiently large t without reference to $\overline{\mathscr{S}}(\mu_{\le t'})$ for $t' > t$.

Theorem 9. *If each agent eventually prefers constant assignments, then there exists a perfect* $\overline{\alpha}$*-stable matching.*

9 Summary and Conclusion

Agents often must act together to generate surplus and the details of the interaction leading to an outcome are very complex. The theory developed above offers a flexible framework to analyze markets where, in addition, time plays a critical role. The template is simple and adaptable. An outcome $\mu = (\mu_1, \ldots, \mu_T)$ is unstable if a coalition has a feasible joint action σ_t

in some period t such that for each coalition member i,

$$\big(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)\big) \succ_i \mu(i) \quad \text{for all} \quad \sigma_{>t} \in \underbrace{\bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|i)}}_{\substack{\text{set of period-}t\text{ outcomes agent } i \text{ thinks} \\ \text{may arise if the coalition pursues } \sigma_t}} \overbrace{\mathscr{S}\big((\mu_{<t}, \tilde{\sigma}_t)\big)}^{\substack{\text{set of future outcomes that} \\ \text{are stable given } (\mu_{<t}, \tilde{\sigma}_t)}}.$$

The sets "$\mathscr{A}(\cdot|\cdot)$" and "$\mathscr{S}(\cdot)$" parameterize a family of solutions by adjusting agents' beliefs about contemporaneous and future market developments, respectively. In applications, these sets' definitions may draw on the market's institutional details or incorporate behavioral refinements. Sharper predictions are likely to result. The proposals above, however, are natural departure points for an initial analysis. And, the application of the proposed framework to problems beyond two-sided matching is promising as well.

A Appendix: Proofs

Proof of Theorem 1. It is sufficient to show that $\mathscr{S}(\mu_{<t}) \neq \varnothing$ for every $\mu_{<t} \in \mathscr{M}_{<t}$ and t. The theorem follows from the $t = 1$ case. Let $\mu_{<T} \in \mathscr{M}_{<T}$. By Lemma 1, $\mu_T \in \mathscr{S}(\mu_{<T})$ if and only if μ_T is a stable matching in a one-period economy where the preference of each agent i, P_i, is given by (2). Such an assignment exists (Gale and Shapley, 1962). Thus, $\mathscr{S}(\mu_{<T}) \neq \varnothing$.

Proceeding by induction, suppose that for every $t' > t$ and all $\mu_{<t'} \in \mathscr{M}_{<t'}$, $\mathscr{S}(\mu_{<t'}) \neq \varnothing$. Let $\mu_{<t} \in \mathscr{M}_{<t}$. For each $i \in I$, define a ranking P_i of agent i's potential partners as follows:

$$j P_i k \iff \begin{aligned} &\min_{\tilde{\mu}_t \in \mathscr{A}(i,j)} \ \min_{\tilde{\mu}_{>t} \in \mathscr{S}((\mu_{<t}, \tilde{\mu}_t))} \big(\mu_{<t}(i), \tilde{\mu}_t(i), \tilde{\mu}_{>t}(i)\big) \\ &\quad \succ_i \min_{\tilde{\mu}'_t \in \mathscr{A}(i,k)} \ \min_{\tilde{\mu}'_{>t} \in \mathscr{S}((\mu_{<t}, \tilde{\mu}'_t))} \big(\mu_{<t}(i), \tilde{\mu}'_t(i), \tilde{\mu}'_{>t}(i)\big). \end{aligned} \tag{A.1}$$

We interpret the righthand side of (A.1) as follows. For each $\tilde{\mu}_t \in \mathscr{A}$ such that $\tilde{\mu}_t(i) = j$, $\tilde{\mu}_{>t} \in \mathscr{M}_{>t}$ is such that $(\mu_{<t}, \tilde{\mu}_t, \tilde{\mu}_{>t})$ cannot be cautiously α-blocked in any period $t' > t$. By the induction hypothesis, $\mathscr{S}((\mu_{<t}, \tilde{\mu}_t)) \neq \varnothing$; hence, $\tilde{\mu}_{>t}$ exists. We find the least favorable matching of this form with respect to $\succ_i$. This is the entry to the left of "$\succ_i$" in (A.1). To the right of "$\succ_i$," the same steps are repeated but $\tilde{\mu}_t(i) = k$. Thus, $j P_i k$ if and only if the "worst" outcome to i

of matching with j is better than the "worst" outcome of matching with k.[20]

Let $\widehat{\mu}_t \in \mathscr{A}$ be a stable assignment in a one-period market where the preference of each agent i is P_i, as defined in (A.1). Such an assignment exists (Gale and Shapley, 1962). By the induction hypothesis, $\mathscr{S}((\mu_{<t}, \widehat{\mu}_t)) \neq \varnothing$. Select some $\widehat{\mu}_{>t} \in \mathscr{S}((\mu_{<t}, \widehat{\mu}_t))$. We argue that $\widehat{\mu}_{\geq t} = (\widehat{\mu}_t, \widehat{\mu}_{>t})$ belongs to $\mathscr{S}(\mu_{<t})$. Suppose the contrary. Since $\widehat{\mu}_{>t} \in \mathscr{S}((\mu_{<t}, \widehat{\mu}_t))$, $(\mu_{<t}, \widehat{\mu}_{\geq t})$ must be blocked in period t. Thus, there exist $K \subseteq I$ and $\sigma_t \in \mathscr{A}(K)$ such that for all $i \in K$,

$$\big(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)\big) \succ_i \big(\mu_{<t}(i), \widehat{\mu}_{\geq t}(i)\big) \quad \text{for all} \quad \sigma_{>t} \in \bigcup_{\widetilde{\sigma}_t \in \mathscr{A}(\sigma_t|i)} \mathscr{S}\big((\mu_{<t}, \widetilde{\sigma}_t)\big). \tag{A.2}$$

There are two cases.

Case 1. Suppose $\sigma_t(i) = i \neq \widehat{\mu}_t(i)$ for some $i \in K$. In this case, $\mathscr{A}(\sigma_t|i) = \mathscr{A}(i,i)$ and (A.2) is equivalent to $\min_{\widetilde{\sigma}_t \in \mathscr{A}(i,i)} \min_{\widetilde{\sigma}_{>t} \in \mathscr{S}((\mu_{<t}, \widetilde{\sigma}_t))} (\mu_{<t}(i), \widetilde{\sigma}_t(i), \widetilde{\sigma}_{>t}(i)) \succ_i (\mu_{<t}(i), \widehat{\mu}_t(i), \widehat{\mu}_{>t}(i))$. As above, the minimizations are with respect to the order $\succ_i$. Hence,

$$\begin{aligned} &\min_{\widetilde{\sigma}_t \in \mathscr{A}(i,i)} \min_{\widetilde{\sigma}_{>t} \in \mathscr{S}((\mu_{<t}, \widetilde{\sigma}_t))} \big(\mu_{<t}(i), \widetilde{\sigma}_t(i), \widetilde{\sigma}_{>t}(i)\big) \\ &\qquad \succ_i \min_{\widetilde{\mu}_t \in \mathscr{A}(i, \widehat{\mu}_t(i))} \min_{\widetilde{\mu}_{>t} \in \mathscr{S}((\mu_{<t}, \widetilde{\mu}_t))} \big(\mu_{<t}(i), \widetilde{\mu}_t(i), \widetilde{\mu}_{>t}(i)\big). \end{aligned}$$

But this implies $i P_i \widehat{\mu}_t(i)$. Thus, $\widehat{\mu}_t$ cannot be a stable assignment in a one-period economy where P_i is the preference of agent i.

Case 2. Suppose $\sigma_t(i) = j \neq \widehat{\mu}(i)$, for some $i \in K$. Thus, $j \in K$ and $\sigma_t(j) = i$. Noting that $\mathscr{A}(i,j) = \mathscr{A}(\sigma_t|i)$ and $\mathscr{A}(j,i) = \mathscr{A}(\sigma_t|j)$, a parallel argument to that of case 1 shows that $j P_i \widehat{\mu}_t(i)$ and $i P_j \widehat{\mu}_t(j)$. Thus, $\widehat{\mu}_t$ cannot be stable. Agents i and j with preferences P_i and P_j, respectively, are a blocking coalition.

As each possible case leads to a contradiction, we conclude that $\widehat{\mu}_{\geq t} \in \mathscr{S}(\mu_{<t})$. □

Proof of Theorem 2. It suffices to show that $\mathscr{C}(\mu_{<t}) \subseteq \mathscr{S}(\mu_{<t})$ for every $\mu_{<t} \in \mathscr{M}_{<t}$ and t. Let $\mu_{<T} \in \mathscr{M}_{<T}$. The assignment μ_T is in $\mathscr{C}(\mu_{<T})$ if and only if it is a core matching in a one-period economy where the preference of each agent i, P_i, is given by (2). That economy's core is not empty and equals the pairwise stable set (Gale and Shapley, 1962). Thus, $\mathscr{C}(\mu_{<T}) = \mathscr{S}(\mu_{<T})$.

[20] Sasaki and Toda (1996) also define a preference ranking for each agent by homing in on worst case outcomes. Our definition differs since we additionally exploit our model's temporal structure—$\widetilde{\mu}_{>t}$ must belong to $\mathscr{S}((\mu_{<t}, \widetilde{\mu}_t))$. The unconstrained worst case selects $\widetilde{\mu}_{>t}$ from $\mathscr{M}_{>t}$.

Proceeding by induction, suppose $\mathscr{C}(\mu_{<t'}) \subseteq \mathscr{S}(\mu_{<t'})$ for all $\mu_{<t'} \in \mathscr{M}_{<t'}$ and all $t' > t$. Fix $\mu_{<t} \in \mathscr{M}_{<t}$. Suppose $\mu_{\geq t} \in \mathscr{C}(\mu_{<t})$. If $\mu_{\geq t} \notin \mathscr{S}(\mu_{<t})$, then $(\mu_{<t}, \mu_{\geq t})$ must be cautiously α-blocked in some period $t' \geq t$. Since $\mathscr{C}(\mu_{<t'}) \subseteq \mathscr{S}(\mu_{<t'})$ for all $t' > t$, $(\mu_{<t}, \mu_{\geq t})$ must be cautiously α-blocked in period t. Thus, there exist $K \subseteq I$ and $\sigma_t \in \mathscr{A}(K)$ such that for each $i \in K$, $(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)) \succ_i \mu(i)$ for all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|i)} \mathscr{S}((\mu_{<t}, \tilde{\sigma}_t))$. Recall that $\mathscr{A}(\sigma_t|K) \subseteq \mathscr{A}(\sigma_t|i)$ if $i \in K$. And, by the induction hypothesis, $\mathscr{C}((\mu_{<t}, \tilde{\sigma}_t)) \subseteq \mathscr{S}((\mu_{<t}, \tilde{\sigma}_t))$ for all $\tilde{\sigma}_t \in \mathscr{A}$. Therefore, coalition K can α-block $(\mu_{<t}, \mu_{\geq t})$ in period t—a contradiction. Thus, the assumption that $\mu_{\geq t} \notin \mathscr{S}(\mu_{<t})$ was incorrect. And so, $\mathscr{C}(\mu_{<t}) \subseteq \mathscr{S}(\mu_{<t})$. □

The next lemma is invoked in the proof of Theorem 3. It is an immediate consequence of history independence and its proof is omitted.

Lemma A.1. *Suppose each agent's preference is history independent. Let $\mu_{<t}, \mu'_{<t} \in \mathscr{M}_{<t}$. (a) $\mathscr{S}(\mu_{<t}) = \mathscr{S}(\mu'_{<t})$. (b) $\mathscr{C}(\mu_{<t}) = \mathscr{C}(\mu'_{<t})$.*

Proof of Theorem 3. It is sufficient to show that $\mathscr{C}(\mu_{<t}) = \mathscr{S}(\mu_{<t})$ for every $\mu_{<t} \in \mathscr{M}_{<t}$ and t. The proof of Theorem 2 shows that $\mathscr{C}(\mu_{<T}) = \mathscr{S}(\mu_{<T})$ for all $\mu_{<T} \in \mathscr{M}_{<T}$. Proceeding by induction, suppose $\mathscr{C}(\mu_{<t'}) = \mathscr{S}(\mu_{<t'})$ for all $\mu_{<t'}$ and $t' > t$. Let $\mu_{<t} \in \mathscr{M}_{<t}$. Since $\mathscr{C}(\mu_{<t}) \subseteq \mathscr{S}(\mu_{<t})$ (see proof of Theorem 2), it is sufficient to establish the converse inclusion. Accordingly, let $\mu_{\geq t} \in \mathscr{S}(\mu_{<t})$. Suppose $(\mu_{<t}, \mu_{\geq t})$ can be α-blocked in period $t' \geq t$. If $t' > t$, then $\mu_{\geq t'} \in \mathscr{S}((\mu_{<t}, \mu_t, \ldots, \mu_{t'-1})) = \mathscr{C}((\mu_{<t}, \mu_t, \ldots, \mu_{t'-1}))$. The equality follows from the induction hypothesis. And so, $(\mu_{<t}, \mu_{\geq t})$ cannot be α-blocked in period $t' > t$.

Therefore, $(\mu_{<t}, \mu_{\geq t})$ must be α-blocked in period t. Thus, there exist $K \subseteq I$ and $\sigma_t \in \mathscr{A}(K)$ such that for each $i \in K$, $(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)) \succ_i (\mu_{<t}(i), \mu_{\geq t}(i))$ for all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|K)} \mathscr{C}((\mu_{<t}, \tilde{\sigma}_t))$. By the induction hypothesis, $\mathscr{C}((\mu_{<t}, \tilde{\sigma}_t)) = \mathscr{S}((\mu_{<t}, \tilde{\sigma}_t))$ for each $\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|K)$. Furthermore, by Lemma A.1, $\mathscr{S}((\mu_{<t}, \tilde{\sigma}_t)) = \mathscr{S}((\mu_{<t}, \sigma'_t))$ where $\sigma'_t \in \mathscr{A}$ is arbitrary. Thus, for each $i \in K$ and all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|i)} \mathscr{S}((\mu_{<t}, \tilde{\sigma}_t))$, $(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)) \succ_i (\mu_{<t}(i), \mu_{\geq t}(i))$. But this implies $\mu_{\geq t} \notin \mathscr{S}(\mu_{<t})$—a contradiction. And so, $\mu_{\geq t} \in \mathscr{C}(\mu_{<t})$. □

Proof of Theorem 4. It is sufficient to show that $\mathscr{C}(\mu_{<t}) \neq \varnothing$ for every $\mu_{<t} \in \mathscr{M}_{<t}$ and t. The proof of Theorem 2 shows that $\mathscr{C}(\mu_{<T}) = \mathscr{S}(\mu_{<T}) \neq \varnothing$ for all $\mu_{<T} \in \mathscr{M}_{<T}$. Proceeding by induction, suppose $\mathscr{C}(\mu_{<t'}) \neq \varnothing$ for all $\mu_{<t'} \in \mathscr{M}_{<t'}$ and $t' > t$. Let $\mu_{<t} \in \mathscr{M}_{<t}$. For each $i \in I$, define a ranking P_i of agent i's potential partners as follows:

$$j P_i k \iff \big(\mu_{<t}(i), j, \ldots, j\big) \succ_i \big(\mu_{<t}(i), k, \ldots, k\big). \tag{A.3}$$

When $t = 1$ in (A.3), P_i reduces to the "isolated preference relation" of Kennes et al. (2014) or the "ex ante spot ranking induced by $\succ_i$" of Kadam and Kotowski (2018a). Let $\mu_{\geq t} = (\mu_t, \ldots, \mu_T)$ be a matching where each μ_s, $s = t, \ldots, T$, is the same core assignment from a one-period market where P_i is the preference of each agent i. Such an assignment exists (Gale and Shapley, 1962). Next, we argue that $\mu_{\geq t} \in \mathscr{C}(\mu_{<t})$. Suppose the contrary. Thus, there exist $K \subseteq I$, $s \geq t$, and $\sigma_s \in \mathscr{A}(K)$, such that $\sigma_s(j) \neq \mu_s(j)$ for some $j \in K$ and for each $i \in K$,

$$\big(\mu_{<t}(i), \mu_t(i), \ldots, \mu_{s-1}(i), \sigma_s(i), \sigma_{>s}(i)\big) \succ_i \big(\mu_{<t}(i), \mu_{\geq t}(i)\big)$$

for all $\sigma_{>s} \in \bigcup_{\tilde{\sigma}_s \in \mathscr{A}(\sigma_s|K)} \mathscr{C}((\mu_{<t}, \mu_t, \ldots, \mu_{s-1}, \tilde{\sigma}_s))$. (Recall that $\mathscr{C}((\mu_{<t}, \mu_t, \ldots, \mu_{s-1}, \tilde{\sigma}_s)) \neq \varnothing$ by the induction hypothesis.) There are two possible cases.

Case 1. Let $i \in K$ and suppose $\sigma_s(i) = i \neq \mu_s(i)$. Since $\mu_t(i) = \cdots = \mu_T(i)$, SSA implies that $(\mu_{<t}(i), i, \ldots, i) \succ_i (\mu_{<t}(i), \mu_t(i), \ldots, \mu_T(i))$. Hence, $i P_i \mu_t(i)$. But this contradicts μ_t being in the core of a one-period market.

Case 2. Let $i \in K$ and suppose $\sigma_s(i) = j \neq \mu_s(i)$. Then $j \in K$ and $\sigma_s(j) = i \neq \mu_{t'}(j)$. The same reasoning as in case 1 lets us conclude that $j P_i \mu_t(i)$ and $i P_j \mu_t(j)$. Thus, μ_t cannot be a core assignment in the corresponding one-period market. Agents i and j with preferences P_i and P_j, respectively, are a blocking coalition.

As each case leads to a contradiction, we conclude that $\mu_{\geq t} \in \mathscr{C}(\mu_{<t})$. □

The next lemma is used in the proof of Theorem 5.

Lemma A.2. *Fix $i \in I$ and $\mu_{<t} \in \mathscr{M}_{<t}$. Suppose $\mathscr{S}^*((\mu_{<t}, \mu_t)) \neq \varnothing$ for all $\mu_t \in \mathscr{A}$.*

(a) Suppose $\sigma_t, \sigma'_t \in \mathscr{A}$ and $\sigma_t(i) = \sigma'_t(i)$, then $\mathscr{A}^*_{\mu_{<t}}(\sigma_t|i) = \mathscr{A}^*_{\mu_{<t}}(\sigma'_t|i)$.

(b) The set $\mathscr{A}^*_{\mu_{<t}}(\sigma_t|i)$ is not empty for every $\sigma_t \in \mathscr{A}$.

Proof of Lemma A.2. Part (a) follows from Definition 5 and the fact that $\mathscr{A}(\sigma_t|i) = \mathscr{A}(\sigma'_t|i)$ whenever $\sigma_t(i) = \sigma'_t(i)$. To prove part (b), fix $i \in I$, $\mu_{<t} \in \mathscr{M}_{<t}$, and $\sigma_t \in \mathscr{A}$. Let $J = I \setminus \{i, \sigma_t(i)\}$. If $J = \varnothing$, then by Definition 5, no assignment in $\mathscr{A}(\sigma_t|i)$ can be considered cautiously α-dominated. Similarly, if $J = \{j\}$, then $\tilde{\sigma}_t(j) = j$ for all $\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|i)$.[21] Thus, there is no alternative assignment among agents in J that agent j strictly prefers. Hence, no assignment in $\mathscr{A}(\sigma_t|i)$ is cautiously α-dominated. In the preceding cases, $\mathscr{A}^*_{\mu_{<t}}(\sigma_t|i) = \mathscr{A}(\sigma_t|i) \neq \varnothing$.

[21] There are two possibilities: (i) $I = \{i, j\}$ and $\sigma_t(i) = i$, which necessarily means $\sigma_t(j) = j$. In this case $\mathscr{A}(i|\sigma_t) = \{\sigma_t\}$. (ii) $I = \{i, k, j\}$ and $\sigma_t(i) = k$, which necessarily means $\sigma_t(j) = j$. Again, $\mathscr{A}(i|\sigma_t) = \{\sigma_t\}$.

Henceforth, suppose $|J| \geq 2$. For each $j \in J$ define a ranking P_j of potential partners in J (including him/herself) as follows:

$$kP_j\ell \iff \begin{aligned} &\min_{\widetilde{\mu}_t \in \mathscr{A}(\sigma_t|i)\cap\mathscr{A}(j,k)} \min_{\widetilde{\mu}_{>t} \in \mathscr{S}^*((\mu_{<t},\widetilde{\mu}_t))} \big(\mu_{<t}(j), k, \widetilde{\mu}_{>t}(j)\big) \\ &\quad \succ_j \min_{\widetilde{\mu}_t \in \mathscr{A}(\sigma_t|i)\cap\mathscr{A}(j,\ell)} \min_{\widetilde{\mu}_{>t} \in \mathscr{S}^*((\mu_{<t},\widetilde{\mu}_t))} \big(\mu_{<t}(j), \ell, \widetilde{\mu}_{>t}(j)\big). \end{aligned} \tag{A.4}$$

In (A.4), $\mathscr{A}(\sigma_t|i)\cap\mathscr{A}(j,k)$ is the set of one-period assignments where i is assigned to $\sigma_t(i)$ and j is assigned to k. Since neither j nor k can equal i or $\sigma_t(i)$, this set is not empty. By assumption, $\mathscr{S}^*((\mu_{<t},\widetilde{\mu}_t)) \neq \varnothing$ for all $\widetilde{\mu}_t$. Thus, P_j is well-defined for each $j \in J$.

Define $\nu_t \in \mathscr{A}$ as follows. For i and $\sigma_t(i)$, let $\nu_t(i) = \sigma_t(i)$ and $\nu_t(\sigma_t(i)) = i$. On J, let $\nu_t(\cdot)$ be any stable assignment among agents in J where the preference of each $j \in J$ is P_j, as defined in (A.4). Such an assignment exists (Gale and Shapley, 1962). Observe that $\nu_t \in \mathscr{A}(\sigma_t|i)$.

We argue that $\nu_t \in \mathscr{A}^*_{\mu_{<t}}(\sigma_t|i)$. Suppose the contrary. Thus, there exist $K \subseteq J$ and $\widehat{\nu}_t \in \mathscr{A}(\nu_t|i)\cap\mathscr{A}(K)$ such that for all $j \in K$,

$$\begin{aligned} &\min_{\widehat{\nu}'_t \in \mathscr{A}(\widehat{\nu}_t|j)} \min_{\widehat{\nu}'_{>t} \in \mathscr{S}^*((\mu_{<t},\widehat{\nu}'_t))} \big(\mu_{<t}(j), \widehat{\nu}_t(j), \widehat{\nu}'_{>t}(j)\big) \\ &\quad \succ_j \max_{\nu'_t \in \mathscr{A}(\nu_t|j)} \max_{\nu'_{>t} \in \mathscr{S}^*((\mu_{<t},\widehat{\nu}'_t))} \big(\mu_{<t}(j), \nu_t(j), \nu'_{>t}(j)\big). \end{aligned} \tag{A.5}$$

In (A.5), which follows from Definition 5, the minimizations and maximizations are with respect to the preference $\succ_j$. But this implies

$$\begin{aligned} &\min_{\widehat{\nu}'_t \in \mathscr{A}(\widehat{\nu}_t|j)} \min_{\widehat{\nu}'_{>t} \in \mathscr{S}^*((\mu_{<t},\widehat{\nu}'_t))} \big(\mu_{<t}(j), \widehat{\nu}_t(j), \widehat{\nu}'_{>t}(j)\big) \\ &\quad \succ_j \min_{\nu'_t \in \mathscr{A}(\nu_t|j)} \min_{\nu'_{>t} \in \mathscr{S}^*((\mu_{<t},\widehat{\nu}'_t))} \big(\mu_{<t}(j), \nu_t(j), \nu'_{>t}(j)\big). \end{aligned} \tag{A.6}$$

Fix $j \in K \subseteq J$. If $\widehat{\nu}_t(j) = k$, then (A.6) implies that $kP_j\,\nu_t(j)$. (The case where $j = k$ is possible.) Since $\widehat{\nu}_t \in \mathscr{A}(K)$, $k \in K$ as well. By similar reasoning, $jP_k\,\nu_t(k)$. But, if $kP_j\,\nu_t(j)$ and $jP_k\,\nu_t(k)$, then ν_t is not a stable matching among agents in J, a contradiction. Thus, $\nu_t \in \mathscr{A}^*_{\mu_{<t}}(\sigma_t|i)$. □

Proof of Theorem 5. It is sufficient to show that $\mathscr{S}^*(\mu_{<t}) \neq \varnothing$ for every $\mu_{<t} \in \mathscr{M}_{<t}$ and t. When $t = T$, cautious α^*-blocking and cautious α-blocking coincide. Thus, $\mathscr{S}^*(\mu_{<T}) = \mathscr{S}(\mu_{<T}) \neq \varnothing$ for every $\mu_{<T} \in \mathscr{M}_{<T}$. Proceeding by induction, suppose $\mathscr{S}^*(\mu_{<t'}) \neq \varnothing$ for every $t' > t$ and every $\mu_{<t'} \in \mathscr{M}_{<t'}$. Fix $\mu_{<t} \in \mathscr{M}_{<t}$. For $i, j \in I$ let

$$\mathscr{A}^*_{\mu_{<t}}(i,j) \coloneqq \bigcup_{\mu_t \in \mathscr{A}(i,j)} \mathscr{A}^*_{\mu_{<t}}(\mu_t|i) \tag{A.7}$$

be the set of assignments where i is assigned to j that are not considered cautiously α-dominated by agent i. Since $\mathscr{A}^*_{\mu_{<t}}(\mu_t|i) \neq \varnothing$, $\mathscr{A}^*_{\mu_{<t}}(i,j) \neq \varnothing$. Next, define a ranking P_i^* of potential partners for agent i as follows:

$$j P_i^* k \iff \min_{\widetilde{\mu}_t \in \mathscr{A}^*_{\mu_{<t}}(i,j)} \; \min_{\widetilde{\mu}_{>t} \in \mathscr{S}^*((\mu_{<t},\widetilde{\mu}_t))} \big(\mu_{<t}(i), \widetilde{\mu}_t(i), \widetilde{\mu}_{>t}(i)\big) \succ_i \min_{\widetilde{\mu}'_t \in \mathscr{A}^*_{\mu_{<t}}(i,k)} \; \min_{\widetilde{\mu}'_{>t} \in \mathscr{S}^*((\mu_{<t},\widetilde{\mu}'_t))} \big(\mu_{<t}(i), \widetilde{\mu}'_t(i), \widetilde{\mu}'_{>t}(i)\big).$$

The ranking P_i^* is like P_i from (A.1) except $\mathscr{A}^*_{\mu_{<t}}$ and $\mathscr{S}^*$ replace $\mathscr{A}$ and $\mathscr{S}$, respectively.

Let $\widehat{\mu}_t$ be a stable assignment in a one-period market where the preference of each agent i is P_i^*. Such an assignment exists (Gale and Shapley, 1962). By the induction hypothesis, $\mathscr{S}^*((\mu_{<t},\widehat{\mu}_t)) \neq \varnothing$. Select some $\widehat{\mu}_{>t} \in \mathscr{S}^*((\mu_{<t},\widehat{\mu}_t))$. We argue that $\widehat{\mu}_{\geq t} = (\widehat{\mu}_t, \widehat{\mu}_{>t}) \in \mathscr{S}^*(\mu_{<t})$. Suppose the contrary. Since $\widehat{\mu}_{>t} \in \mathscr{S}^*((\mu_{<t},\widehat{\mu}_t))$, $(\mu_{<t},\widehat{\mu}_{\geq t})$ must be be cautiously α^*-blocked in period t. Thus, there exist $K \subseteq I$ and $\sigma_t \in \mathscr{A}(K)$ such that for each $i \in K$, $(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)) \succ_i (\mu_{<t}(i), \widehat{\mu}_t(i), \widehat{\mu}_{\geq t}(i))$ for all $\sigma_{>t} \in \bigcup_{\widetilde{\sigma}_t \in \mathscr{A}^*_{\mu_{<t}}(\sigma_t|i)} \mathscr{S}^*((\mu_{<t},\widetilde{\sigma}_t))$. There are two cases.

Case 1. Suppose $\sigma_t(i) = i \neq \widehat{\mu}_t(i)$ for some $i \in K$. By Lemma A.2 and the definition of $\mathscr{A}^*_{\mu_{<t}}(\cdot,\cdot)$ in (A.7), $\mathscr{A}^*_{\mu_{<t}}(\sigma_t|i) = \mathscr{A}^*_{\mu_{<t}}(i,i)$. Thus,

$$\min_{\widetilde{\sigma}_t \in \mathscr{A}^*_{\mu_{<t}}(i,i)} \; \min_{\sigma_{>t} \in \mathscr{S}^*((\mu_{<t},\widetilde{\sigma}_t))} \big(\mu_{<t}(i), i, \sigma_{>t}(i)\big) \succ_i \big(\mu_{<t}(i), \widehat{\mu}_t(i), \widehat{\mu}_{>t}(i)\big).$$

And hence,

$$\min_{\widetilde{\sigma}_t \in \mathscr{A}^*_{\mu_{<t}}(i,i)} \; \min_{\sigma_{>t} \in \mathscr{S}^*((\mu_{<t},\widetilde{\sigma}_t))} \big(\mu_{<t}(i), i, \sigma_{>t}(i)\big) \succ_i \min_{\widetilde{\mu}_t \in \mathscr{A}^*_{\mu_{<t}}(i,\widehat{\mu}_t(i))} \; \min_{\widetilde{\mu}_{>t} \in \mathscr{S}^*((\mu_{<t},\widetilde{\mu}_t))} \big(\mu_{<t}(i), \widehat{\mu}_t(i), \widetilde{\mu}_{>t}(i)\big),$$

which implies $i P_i^* \widehat{\mu}_t(i)$. Thus, $\widehat{\mu}_t$ cannot be a stable matching in a one-period economy where the preference of agent i is P_i^*.

Case 2. Suppose $\sigma_t(i) = j \neq \widehat{\mu}(i)$, for some $i \in K$. Thus, $j \in K$ and $\sigma_t(j) = i$. Noting that $\mathscr{A}^*_{\mu_{<t}}(i,j) = \mathscr{A}^*_{\mu_{<t}}(\sigma_t|i)$ and $\mathscr{A}^*_{\mu_{<t}}(j,i) = \mathscr{A}^*_{\mu_{<t}}(\sigma_t|j)$, a parallel argument to that of case 1 shows that $j P_i^* \widehat{\mu}_t(i)$ and $i P_j^* \widehat{\mu}_t(j)$. Thus, $\widehat{\mu}_t$ cannot be stable. Agents i and j with preferences P_i^* and P_j^*, respectively, are a blocking coalition.

As each case leads to a contradiction, we conclude that $\widehat{\mu}_{\geq t} \in \mathscr{S}^*(\mu_{<t})$. □

Proof of Theorem 6. It is sufficient to show that $\mathscr{S}^*(\mu_{<t}) \subseteq \mathscr{S}(\mu_{<t})$ for all $\mu_{<t} \in \mathscr{M}_{<t}$ and t. The proof of Theorem 5 shows that $\mathscr{S}^*(\mu_{<T}) = \mathscr{S}(\mu_{<T})$ for all $\mu_{<T} \in \mathscr{M}_{<T}$. Proceeding by

induction, suppose $\mathscr{S}^*(\mu_{<t'}) \subseteq \mathscr{S}(\mu_{<t'})$ for all $\mu_{<t'} \in \mathscr{M}_{<t'}$ and $t' > t$. Fix $\mu_{\geq t} \in \mathscr{S}^*(\mu_{<t})$. If $\mu_{\geq t} \notin \mathscr{S}(\mu_{<t})$, then $(\mu_{<t}, \mu_{\geq t})$ can be cautiously α-blocked in some period $t' \geq t$. Since $\mathscr{S}^*(\mu_{<t'}) \subseteq \mathscr{S}(\mu_{<t'})$ for all $t' > t$, $(\mu_{<t}, \mu_{\geq t})$ must be cautiously α-blocked in period t. Thus, there exist $K \subseteq I$ and $\sigma_t \in \mathscr{A}(K)$ such that for each $i \in K$, $(\mu_{<t}(i), \sigma_t(i), \sigma_{>t}(i)) \succ_i \mu(i)$ for all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|i)} \mathscr{S}((\mu_{<t}, \tilde{\sigma}_t))$. Recall that $\mathscr{A}^*_{\mu_{<t}}(\sigma_t|i) \subseteq \mathscr{A}(\sigma_t|i)$. And, by the induction hypothesis, $\mathscr{S}^*((\mu_{<t}, \tilde{\sigma}_t)) \subseteq \mathscr{S}((\mu_{<t}, \tilde{\sigma}_t))$ for all $\tilde{\sigma}_t \in \mathscr{A}$. Therefore, coalition K can cautiously α^*-block $(\mu_{<t}, \mu_{\geq t})$ in period t—a contradiction. Thus, the assumption that $\mu_{\geq t} \notin \mathscr{S}(\mu_{<t})$ is incorrect. And so, $\mathscr{S}^*(\mu_{<t}) \subseteq \mathscr{S}(\mu_{<t})$. □

Proof of Theorem 7. Since preferences are time-separable, the proof follows from known results and we only sketch the argument. We rely on Shapley and Shubik (1971) and Roth and Sotomayor (1990, Chapter 7) for results on the one-period model.

For every $\gamma_{<T}$, $\mathscr{S}(\gamma_{<T}) = \mathscr{C}(\gamma_{<T}) \neq \varnothing$. This is because our solutions coincide with the standard definitions of stability and the core in the final period. Shapley and Shubik (1971) show that the stable set and the core are nonempty and equal.

Given that $\mathscr{S}(\gamma_{<T}) = \mathscr{C}(\gamma_{<T}) \neq \varnothing$ for all $\gamma_{<T}$, the conclusions of Theorem 2, Lemma A.1, and Theorem 3 apply to this model. The latter two results are true because agents' preferences are history independent. The proofs of these results are essentially identical to those presented above and are omitted. Thus, for all $\gamma_{<t}$, $\mathscr{C}(\gamma_{<t}) = \mathscr{S}(\gamma_{<t})$. Therefore, to prove the theorem it is sufficient to show that $\mathscr{S}(\gamma_{<t}) \neq \varnothing$.

Proceeding by induction, suppose $\mathscr{S}(\gamma_{<t'}) \neq \varnothing$ for all outcomes $\gamma_{<t'}$ where $t' > t$. Fix $\gamma_{<t}$ and consider the period-t market as an instance of Shapley and Shubik's (1971) model. Each agent i strictly prefers the outcome $\gamma_t = (x_t, u_t)$ over $\gamma'_t = (x'_t, u'_t)$ if and only if $u^i_t > u'^i_t$. Let $\hat{\gamma}_t$ be any stable (equivalently, core) outcome in this one-period economy. Choose any $\hat{\gamma}_{>t} \in \mathscr{S}((\gamma_{<t}, \hat{\gamma}_t))$. We will verify that $\hat{\gamma}_{\geq t} = (\hat{\gamma}_t, \hat{\gamma}_{>t}) \in \mathscr{S}(\gamma_{<t})$. To do so, it is sufficient to show that $(\gamma_{<t}, \hat{\gamma}_{\geq t})$ cannot be cautiously α-blocked in period t.

Suppose the contrary. Let $(u_{<t}, \hat{u}_{\geq t})$ be the tuple of utilities given outcome $(\gamma_{<t}, \hat{\gamma}_{\geq t})$. Assume coalition K can cautiously α-block $(\gamma_{<t}, \hat{\gamma}_{\geq t})$ in period t with $\sigma_t = (y_t, s_t)$. Thus,

$$\min_{\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t|i)} \mathscr{S}((\gamma_{\leq t}, \tilde{\sigma}_t))} U_i\left(u^i_{<t}, s^i_t, s^i_{>t}\right) > U_i\left(u^i_{<t}, \hat{u}^i_t, \hat{u}^i_{>t}\right)$$

for each $i \in K$. In the preceding expression $s^i_{>t}$ is agent i's utility vector associated with continuation $\sigma_{>t}$. This implies, $U_i(u^i_{<t}, s^i_t, \hat{u}^i_{>t}) > U_i(u^i_{<t}, \hat{u}^i_t, \hat{u}^i_{>t})$ and hence $s^i_t > \hat{u}^i_t$. Thus, coalition K can block $\hat{\gamma}_t$ in a one-period economy—a contradiction. □

Proof of Theorem 8. (a) The proof follows that of Lemma 5.5 in Roth and Sotomayor (1990).

Suppose coalition K can cautiously α-block μ in period t with $\sigma_t \in \mathscr{A}_t$. If there exists a $s \in K \cap S$ such that $\sigma_t(s) = s$ or a $c \in K \cap C$ such that $\sigma_t(c) \subsetneq \mu_t(c)$, then this agent can cautiously α-block μ in period t alone with $\sigma_t \in \mathscr{A}_t$.

Instead, suppose $\sigma_t(c) \nsubseteq \mu_t(c)$ for some $c \in K \cap C$. Thus, $(\mu_{<t}(c), \sigma_t(c), \sigma_{>t}(c)) \succ_c (\mu_{<t}(c), \mu_t(c), \mu_{>t}(c))$ for all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}_t(\sigma_t|c)} \mathscr{S}((\mu_{\leq t}, \tilde{\sigma}_t))$. As the college's preference is conditionally responsive, there exists a strict and responsive ranking $P_c^{\mu_{<t}}$ of 2^S such that $\sigma_t(c) P_c^{\mu_{<t}} \mu_t(c)$. Hence, there exist $s \in \sigma_t(c) \setminus \mu_t(c)$ and $s' \in \mu_t(c) \setminus \sigma_t(c)$ such that $\{s\} P_c^{\mu_{<t}} \{s'\}$ (Roth and Sotomayor, 1990, p. 130). Necessarily, $s \in K \cap S$.

Now consider an assignment $\sigma'_t \in \mathscr{A}_t$ where $\sigma'_t(s) = c$ and $\sigma'_t(c) = \mu_t(c) \cup \{s\} \setminus \{s'\}$. Since $s \in K \cap S$ and $\mathscr{A}_t(\sigma_t|s) = \mathscr{A}_t(\sigma'_t|s)$, $(\mu_{<t}(s), \sigma'_t(s), \sigma_{>t}(s)) \succ_s (\mu_{<t}(s), \mu_t(s), \mu_{>t}(s))$ for all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}_t(\sigma'_t|s)} \mathscr{S}((\mu_{\leq t}, \tilde{\sigma}_t))$. Moreover, since the college's preference is conditionally responsive, $\sigma'_t(c) P_c^{\mu_{<t}} \mu_t(c)$ implies that $(\mu_{<t}(c), \sigma'_t(c), x_{>t}) \succ_c (\mu_{<t}(c), \mu_t(c), y_{>t})$ for all $x_{>t}$ and $y_{>t}$. In particular, $(\mu_{<t}(c), \sigma'_t(c), \sigma_{>t}(c)) \succ_c (\mu_{<t}(c), \mu_t(c), \mu_{>t}(c))$ for all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}_t(\sigma'_t|c)} \mathscr{S}((\mu_{\leq t}, \tilde{\sigma}_t))$. Thus, c and s can cautiously α-block μ in period t.

(b) The argument mirrors the proof of Theorem 1. It is sufficient to show that $\mathscr{S}(\mu_{<t}) \neq \emptyset$ for every $\mu_{<t} \in \mathscr{M}_{<t}$ and t. Let $\mu_{<T} \in \mathscr{M}_{<T}$. Paralleling Lemma 1, $\mu_T \in \mathscr{S}(\mu_{<T})$ if and only if μ_T is group-stable in a one-period economy where the preference P_s of each $s \in S$ satisfies $c P_s c' \iff (\mu_{<T}(s), c) \succ_s (\mu_{<T}(s), c')$ and the preference P_c of each $c \in C$ satisfies $x_T P_c y_T \iff (\mu_{<T}(c), x_T) \succ_c (\mu_{<T}(c), y_T)$. Since colleges' preferences are conditionally responsive, each P_c is a strict and responsive ranking of 2^S. Thus, there exists a group-stable assignment (Roth and Sotomayor, 1990). Hence, $\mathscr{S}(\mu_{<T}) \neq \emptyset$.

Proceeding by induction, suppose that for every $t' > t$ and every $\mu_{<t'} \in \mathscr{M}_{<t'}$, $\mathscr{S}(\mu_{<t'}) \neq \emptyset$. Let $\mu_{<t} \in \mathscr{M}_{<t}$. For each $s \in S$, define a ranking P_s of $C \cup \{s\}$ as follows:

$$c P_s c' \iff \begin{aligned} &\min_{\tilde{\mu}_t \in \mathscr{A}_t(s,c)} \min_{\tilde{\mu}_{>t} \in \mathscr{S}((\mu_{<t}, \tilde{\mu}_t))} \left(\mu_{<t}(s), \tilde{\mu}_t(s), \tilde{\mu}_{>t}(s)\right) \\ &\succ_s \min_{\tilde{\mu}'_t \in \mathscr{A}_t(s,c')} \min_{\tilde{\mu}'_{>t} \in \mathscr{S}((\mu_{<t}, \tilde{\mu}'_t))} \left(\mu_{<t}(s), \tilde{\mu}'_t(s), \tilde{\mu}'_{>t}(s)\right). \end{aligned} \tag{A.8}$$

In (A.8), $\mathscr{A}_t(s,c)$ is the set of $\sigma_t \in \mathscr{A}_t$ such that $\sigma_t(s) = c$ and the minimizations are with respect to $\succ_s$. For each $c \in C$, define a ranking P_c of 2^S as $x_t P_c y_t \iff (\mu_{<t}(c), x_t, x_{>t}) \succ_c (\mu_{<t}(c), y_t, y_{>t})$ where $x_{>t}$ and $y_{>t}$ are fixed arbitrary sequences of feasible assignments for college c for the periods $t' > t$. Since colleges' preferences are conditionally responsive, P_c is a strict and responsive ranking of 2^S. Let $\hat{\mu}_t \in \mathscr{A}_t$ be group-stable in the one-period economy where each student's preference is P_s and each college's preference is P_c, as defined above. Such an assignment exists (Roth and Sotomayor, 1990). Next, select some $\hat{\mu}_{>t} \in \mathscr{S}((\mu_{<t}, \hat{\mu}_t))$.

(By the induction hypothesis, $\mathscr{S}((\mu_{<t},\widehat{\mu}_t)) \neq \varnothing$.) We argue that $\widehat{\mu}_{\geq t} = (\widehat{\mu}_t,\widehat{\mu}_{>t}) \in \mathscr{S}(\mu_{<t})$. To derive a contradiction, suppose this is not true. Since $\widehat{\mu}_{>t} \in \mathscr{S}((\mu_{<t},\widehat{\mu}_t))$, $(\mu_{<t},\widehat{\mu}_{\geq t})$ cannot be cautiously α-blocked in any period $t' > t$. Thus, some coalition K can cautiously α-block $(\mu_{<t},\widehat{\mu}_{\geq t})$ in period t with some $\sigma_t \in \mathscr{A}_t$. There are three cases.

Case 1. Suppose $\sigma_t(s) = s$ for some $s \in K \cap S$. The same argument as in case 1 in the proof of Theorem 1 establishes that $sP_s\widehat{\mu}_t(s)$. Thus, student s can block $\widehat{\mu}_t$ alone and $\widehat{\mu}_t$ cannot be group stable, contradicting its definition.

Case 2. Suppose $\sigma_t(c) \subsetneq \mu_t(c)$ for some $c \in K \cap C$. This implies $(\mu_{<t}(c),\sigma_t(c),\sigma_{>t}(c)) \succ_c (\mu_{<t}(c),\widehat{\mu}_t(c),\widehat{\mu}_{>t}(c))$ for all $\sigma_{>t} \in \bigcup_{\widetilde{\sigma}_t \in \mathscr{A}_t(\sigma_t|c)} \mathscr{S}((\mu_{\leq t},\widetilde{\sigma}_t))$. As the college's preference is conditionally responsive $\sigma_t(c)P_c\widehat{\mu}_t(c)$. Hence, college c can block $\widehat{\mu}_t$ alone and $\widehat{\mu}_t$ cannot be group stable, contradicting its definition.

Case 3. Suppose $\sigma_t(c) \nsubseteq \widehat{\mu}_t(c)$ for some $c \in K \cap C$. As in the proof of part (a) above, there exist $s \in \sigma_t(c) \setminus \widehat{\mu}_t(c)$ and $s' \in \widehat{\mu}_t(c) \setminus \sigma_t(c)$ such that $\{s\}P_c\{s'\}$ and $s \in K \cap S$. Consider $\sigma'_t \in \mathscr{A}_t$ where $\sigma'_t(s) = c$ and $\sigma'_t(c) = \widehat{\mu}_t(c) \cup \{s\} \setminus \{s'\}$. Note that $\sigma'_t(c)P_c\widehat{\mu}_t(c)$. Moreover, $\mathscr{A}_t(s,c) = \mathscr{A}_t(\sigma_t|s) = \mathscr{A}_t(\sigma'_t|s)$. And, as in case 2 of the proof of Theorem 1, we can conclude that $cP_s\widehat{\mu}_t(c)$. But this means $\widehat{\mu}_t$ cannot be group stable, a contradiction.

As each possible case leads to a contradiction, we conclude that $\widehat{\mu}_{\geq t} \in \mathscr{S}(\mu_{<t})$. ☐

Proof of Theorem 9. Let $T \geq \max_i\{T_i\}$ and define the mapping $\overline{\mathscr{S}}(\cdot)$ as follows.

(a) If $t \geq T$ and $\mu_{<t} \in \mathscr{M}_{<t}$, then $\mu_{\geq t} \in \overline{\mathscr{S}}(\mu_{<t})$ if and only if $\mu_{\geq t} = (\mu_t,\mu_{t+1},\ldots)$ is a continuation matching such that $\mu_t = \mu_{t+1} = \cdots$ and $(\mu_{<t}(i),\mu_t(i),\mu_{t+1}(i),\ldots) \succeq_i (\mu_{<t}(i),i,i,\ldots)$ for all i.

(b) If $t < T$ and $\mu_{<t} \in \mathscr{M}_{<t}$, then $\overline{\mathscr{S}}(\mu_{<t})$ is defined via backward induction mimicking the argument in the proof of Theorem 1. That is, $\overline{\mathscr{S}}(\mu_{<t})$ is the set of assignments $\mu_{\geq t}$ such that $(\mu_{<t},\mu_{\geq t})$ cannot be cautiously $\overline{\alpha}$-blocked by any coalition in any period $s \geq t$. The base case for the induction, $\overline{\mathscr{S}}(\mu_{<T})$, is defined in part (a) above.

To prove the theorem it is sufficient to show the following:

(i) For each $t \geq T$, $\mu_{<t} \in \mathscr{M}_{<t}$, and $\mu_{\geq t} \in \overline{\mathscr{S}}(\mu_{<t})$, the matching $\mu = (\mu_{<t},\mu_{\geq t})$ cannot be cautiously $\overline{\alpha}$-blocked by any coalition in any period $s \geq t$.

(ii) For each $t \geq T$ and $\mu_{<t} \in \mathscr{M}_{<t}$, and $\mu_{\geq t} \notin \overline{\mathscr{S}}(\mu_{<t})$, the matching $\mu = (\mu_{<t},\mu_{\geq t})$ can be cautiously $\overline{\alpha}$-blocked by some coalition in some period $s \geq t$.

When $t < T$, the argument in the proof of Theorem 1 applies and completes the proof.

Proof of (i). Fix $t \geq T$ and $\mu_{\geq t} \in \overline{\mathscr{S}}(\mu_{<t})$. To derive a contradiction, suppose $\mu = (\mu_{<t}, \mu_{\geq t})$ can be cautiously $\overline{\alpha}$-blocked in period $s \geq t$ by coalition K with $\sigma_s \in \mathscr{A}(K)$. Suppose $\sigma_s(i) = j \neq \mu_s(i)$ for some $i \in K$ (the case where $j = i$ is possible). Thus,

$$(\mu_{<t}(i), \mu_t(i), \ldots, \mu_{s-1}(i), \underset{\substack{\uparrow \\ \sigma_s(i)}}{j}, \underbrace{k, k, \ldots}_{\sigma_{>s}(i)}) \succ_i (\mu_{<t}(i), \underbrace{\mu_t(i), \ldots, \mu_{s-1}(i), \mu_s(i), \ldots}_{(\mu_t(i) = \mu_{t+1}(i) = \cdots)}) \tag{A.9}$$

for all $\sigma_{>s}(i) = (k, k, \ldots) \in \bigcup_{\tilde{\sigma}_s \in \mathscr{A}(\sigma_s | i)} \overline{\mathscr{S}}((\mu_{<s}, \tilde{\sigma}_s))$. Since agent i eventually prefers constant assignments and $t \geq T \geq T_i$, (A.9) implies that

$$(\mu_{<t}(i), i, i, \ldots) \succeq_i (\mu_{<t}(i), \mu_t(i), \ldots, \mu_{s-1}(i), j, k, k, \ldots) \succ_i (\mu_{<t}(i), \mu_t(i), \mu_{t+1}(i), \ldots). \tag{A.10}$$

However, this contradicts the definition of $\mu_{\geq t}(i) = (\mu_t(i), \mu_{t+1}(i), \ldots)$ in (a) above.

Proof of (ii). Fix $t \geq T$ and $\mu_{\geq t} \notin \overline{\mathscr{S}}(\mu_{<t})$. One of two cases applies.

Case 1. The continuation $\mu_{\geq t}(i)$ is not a constant sequence for some $i \in I$. Given the assumption on preferences, this implies

$$(\mu_{<t}(i), i, i, \ldots) \succ_i (\mu_{<t}(i), \mu_{\geq t}(i)). \tag{A.11}$$

And if $\sigma_t \in \mathscr{A}(i)$, the definition of $\overline{\mathscr{S}}((\mu_{<t}, \sigma_t))$ implies that

$$(\mu_{<t}(i), i, \sigma_{>t}(i)) \succeq_i (\mu_{<t}(i), i, i, i, \ldots) \tag{A.12}$$

for all $\sigma_{>t} \in \bigcup_{\tilde{\sigma}_t \in \mathscr{A}(\sigma_t | i)} \overline{\mathscr{S}}((\mu_{<t}, \tilde{\sigma}_t))$. Together, (A.11) and (A.12) imply that agent i can cautiously $\overline{\alpha}$-block $\mu = (\mu_{<t}, \mu_{\geq t})$ in period t.

Case 2. The continuation $\mu_{\geq t}(i)$ is a constant sequence for each agent i. Thus, from the definition of $\overline{\mathscr{S}}(\mu_{<t})$, there exists some i for whom $(\mu_{<t}(i), i, i, \ldots) \succ_i (\mu_{<t}(i), \mu_{\geq t}(i))$, which is identical to (A.11) above. Noting the preceding argument, (A.12) also holds. Thus, agent i can cautiously $\overline{\alpha}$-block $\mu = (\mu_{<t}, \mu_{\geq t})$ in period t.

Thus, $\mu = (\mu_{<t}, \mu_{\geq t})$ can be cautiously $\overline{\alpha}$-blocked. □

B Appendix: Examples

Example B.1. The strongest solution proposed by Damiano and Lam (2005) is *self-sustaining stability* (S^3). Application of S^3 to our model identifies outcomes that are not perfect α-stable. Consider the following economy with one man and two women:

$$\succ_{m_1}: \underset{\mu}{\mathbf{w_1 w_2}}, \underset{\sigma}{\underline{w_2 w_2}}, m_1 m_1, \ldots \qquad \succ_{w_1}: \underset{\mu}{\mathbf{m_1 w_1}}, \underset{\sigma}{\underline{w_1 w_1}}, \ldots \qquad \succ_{w_2}: \underset{\sigma}{\underline{m_1 m_1}}, \underset{\mu}{\mathbf{w_2 m_1}}, w_2 w_2, \ldots$$

In both μ and σ, each agent receives his/her favorite period-2 assignment. The matching μ must satisfy S^3 since m_1 and w_1 receive their most-preferred outcomes. The matching σ must also be stable in the sense of S^3. Any period-1 blocking coalition must involve everyone. But this is impossible since w_2 becomes worse off. Note that the preceding argument does not rely on the coalition-proofness aspects of S^3. Instead, it depends on blocking coalition members matching only among themselves in all periods.

The only perfect α-stable matching is μ. The matching σ can be α-blocked in period 1 by m_1 and w_1. Thus, perfect α-stability is not weaker than Damiano and Lam's proposal.

Example B.2. Kadam and Kotowski (2018a,b) study a solution they call "dynamic stability." In Example B.1, both μ and σ are dynamically stable. Both matchings are also in the sequential core (see Section 4). Thus, perfect α-stability is not weaker than these solutions.

Example B.3. Kennes et al. (2014) examine the assignment of children to daycares. They propose a mechanism called the DA-IP that identifies a "stable" outcome in their model. To apply their mechanism to our model, we formulate an example that can be embedded in their framework. The example is adapted from Kadam and Kotowski (2018a).

Consider a two-period economy with three men and three women. Their preferences are:

$$\begin{aligned}
&\succ_{m_1}: w_2 w_2, \underset{\mu}{\mathbf{w_1 w_2}}, w_1 w_1, m_1 m_1, w_2 m_1, \ldots && \succ_{w_1}: m_1 m_1, m_2 m_2, m_3 m_3, m_1 m_2, \underset{\mu}{\mathbf{m_1 m_3}}, w_1 w_1, \ldots \\
&\succ_{m_2}: w_1 w_1, \underset{\mu}{\mathbf{w_3 w_3}}, w_3 w_1, m_2 m_2, \ldots && \succ_{w_2}: m_3 m_3, m_1 m_1, \underset{\mu}{\mathbf{m_3 m_1}}, w_2 w_2, m_1 w_2, w_2 w_2, \ldots \\
&\succ_{m_3}: w_1 w_1, \underset{\mu}{\mathbf{w_2 w_1}}, w_2 w_2, m_3 m_3, \ldots && \succ_{w_3}: \underset{\mu}{\mathbf{m_2 m_2}}, w_3 w_3, \ldots
\end{aligned}$$

To nest this economy in the setting of Kennes et al. (2014), interpret the men as the "children" and the women as the "daycares," each with unit capacity. Assume that the agents' priorities at daycares are initially $m_1 \rhd_{w_1} m_2 \rhd_{w_1} m_3$, $m_3 \rhd_{w_2} m_1$, and $m_2 \rhd_{w_3} \cdots$. These priorities correspond to the rankings of the constant assignments in the women's preferences. Priorities for period 2 are derived from the preferences conditional on the period 1 matching. Thus, if m is assigned to w in period 1, his priority at w will be highest among all agents.

The DA-IP mechanism identifies the matching μ. This matching is “stable” in the sense of Kennes et al. (2014, Definition 8) since the preferences and priority structure satisfy their assumptions. The matching μ can be α-blocked in period 1 by m_1 and w_2. Thus, perfect α-stability is not weaker than stability in the sense of Kennes et al. (2014, Definition 8).

Example B.4 (No “Optimal” Stable Matching)**.** Consider a two-period economy with three men and three women. The agents’ preferences are:

$$\succ_{m_1}: \underset{\mu}{\mathbf{w_1 w_2}}, \underset{\nu}{\boxed{w_3 w_1}}, \underset{\sigma}{\underline{w_2 w_3}}, m_1 w_3, m_1 m_1, w_1 m_1, w_3 m_1, w_2 m_1, \ldots$$

$$\succ_{m_2}: \underset{\sigma}{\underline{w_3 w_1}}, \underset{\mu}{\mathbf{w_2 w_3}}, \underset{\nu}{\boxed{w_1 w_2}}, m_2 w_2, m_2 m_2, w_3 m_2, w_2 m_2, w_1 m_2, \ldots$$

$$\succ_{m_3}: \underset{\nu}{\boxed{w_2 w_3}}, \underset{\sigma}{\underline{w_1 w_2}}, \underset{\mu}{\mathbf{w_3 w_1}}, m_3 w_1, m_3 m_3, w_2 m_3, w_1 m_3, w_3 m_3, \ldots$$

$$\succ_{w_1}: \underset{\sigma}{\underline{m_3 m_2}}, \underset{\mu}{\mathbf{m_1 m_3}}, \underset{\nu}{\boxed{m_2 m_1}}, w_1 m_1, w_1 w_1, m_3 w_1, m_1 w_1, m_2 w_1, \ldots$$

$$\succ_{w_2}: \underset{\mu}{\mathbf{m_2 m_1}}, \underset{\nu}{\boxed{m_3 m_2}}, \underset{\sigma}{\underline{m_1 m_3}}, w_2 m_3, w_2 w_2, m_2 w_2, m_3 w_2, m_1 w_2, \ldots$$

$$\succ_{w_3}: \underset{\nu}{\boxed{m_1 m_3}}, \underset{\sigma}{\underline{m_2 m_1}}, \underset{\mu}{\mathbf{m_3 m_2}}, w_3 m_2, w_3 w_3, m_1 w_3, m_2 w_3, m_3 w_3, \ldots$$

Three matchings are highlighted—μ, σ, and ν. Each is perfect α-stable. The men and women all disagree which is best. (In total, there 7 perfect α-stable matchings in this example.)

Example B.5 (Empty Perfect α-Core)**.** The perfect α-core of the economy in Example B.4 is empty. For example, μ can be α-blocked in period 1 by coalition $K = \{m_2, m_3, w_1, w_3\}$ with the assignment $\sigma_1(m_2) = w_3$ and $\sigma_1(m_3) = w_1$. Given this assignment among agents in K, there are two possible assignments among the remaining agents, $I \setminus K = \{m_1, w_2\}$.

(i) Consider $\widetilde{\sigma}_1 \in \mathscr{A}(\sigma_1|K)$ where $\widetilde{\sigma}_1(m_1) = w_2$, $\widetilde{\sigma}_1(m_2) = w_3$, and $\widetilde{\sigma}_1(m_3) = w_1$. Given $\widetilde{\sigma}_1$, there is one period-2 core assignment ($\mathscr{C}(\widetilde{\sigma}_1) = \{\widetilde{\sigma}_2\}$) where $\widetilde{\sigma}_2(m_1) = w_3$, $\widetilde{\sigma}_2(m_2) = w_1$, and $\widetilde{\sigma}_2(m_3) = w_2$. The outcome is identical to σ and all agents in K prefer it over μ.

(ii) Consider $\widetilde{\sigma}'_1 \in \mathscr{A}(\sigma_1|K)$ where $\widetilde{\sigma}'_1(m_1) = m_1$, $\widetilde{\sigma}'_1(m_2) = w_3$, and $\widetilde{\sigma}'_1(m_3) = w_1$. This situation is the same as case (i) except m_1 and w_2 are unmatched in period 1. Given $\widetilde{\sigma}'_1$, there is one core assignment in period 2 ($\mathscr{C}(\widetilde{\sigma}_1') = \{\widetilde{\sigma}_2\}$) where $\widetilde{\sigma}_2(m_1) = w_3$, $\widetilde{\sigma}_2(m_2) = w_1$, and $\widetilde{\sigma}_2(m_3) = w_2$. All agents in K prefer this outcome to μ.

The matchings σ and ν can be blocked similarly. And, the same arguments apply to the three perfect α-stable matchings where only one man and one woman are unmatched in period 1. The matching where all agents are unmatched can be α-blocked by the grand coalition.

References

Ali, S. N. and Liu, C. (2019). Conventions and coalitions in repeated games. Mimeo.

Altınok, A. (2019). Dynamic many-to-one matching. Mimeo.

Ambrus, A. (2006). Coalitional rationalizability. *Quarterly Journal of Economics*, 121(3):903–929.

Aumann, R. J. (1961). The core of a cooperative game without side payments. *Transactions of the American Mathematical Society*, 98(3):539–552.

Aumann, R. J. and Peleg, B. (1960). Von Neumann-Morgenstern solutions to cooperative games without side payments. *Bulletin of the American Mathematical Society*, 66(3):173–179.

Bando, K. (2012). Dynamic matching markets with choice functions. Mimeo.

Becker, R. A. and Chakrabarti, S. K. (1995). The recursive core. *Econometrica*, 63(2):401–423.

Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028.

Bernheim, B. D., Peleg, B., and Whinston, M. D. (1987). Coalition-proof Nash equilibria I. Concepts. *Journal of Economic Theory*, 42(1):1–12.

Corbae, D., Temzelides, T., and Wright, R. (2003). Directed matching and monetary exchange. *Econometrica*, 71(3):731–756.

Damiano, E. and Lam, R. (2005). Stability in dynamic matching markets. *Games and Economic Behavior*, 52(1):34–53.

Diamantoudi, E., Miyagawa, E., and Xue, L. (2015). Decentralized matching: The role of commitment. *Games and Economic Behavior*, 92:1–17.

Doval, L. (2018). A theory of stability in dynamic matching markets. Mimeo.

Doval, L. (2019). Dynamically stable matching. Mimeo.

Du, S. and Livne, Y. (2016). Rigidity of transfers and unravelling in matching markets. Mimeo.

Dur, U. (2012). Dynamic school choice problem. Mimeo.

Gale, D. (1978). The core of a monetary economy without trust. *Journal of Economic Theory*, 19(2):456–491.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.

Habis, H. and Herings, P. J.-J. (2011). Core concepts for incomplete market economies. *Journal of Mathematical Economics*, 47:595–609.

Haeringer, G. and Iehlé, V. (2019). Gradual college admissions. Research Paper 2019-11-03, Baruch College Zicklin School of Business.

Hafalir, I. (2008). Stability of marriage with externalities. *International Journal of Game Theory*, 37:353–369.

Hatfield, J. W. and Kominers, S. D. (2017). Contract design and stability in many-to-many matching. *Games and Economic Behavior*, 101:78–97.

Kadam, S. V. and Kotowski, M. H. (2018a). Multiperiod matching. *International Economic Review*, 59(4):1927–1947.

Kadam, S. V. and Kotowski, M. H. (2018b). Time horizons, lattice structures, and welfare in multi-period matching markets. *Games and Economic Behavior*, 112:1–20.

Kennes, J., Monte, D., and Tumennasan, N. (2014). The daycare assignment: A dynamic matching problem. *American Economic Journal: Microeconomics*, 6(4):362–406.

Kennes, J., Monte, D., and Tumennasan, N. (2019). Strategic performance of deferred acceptance in dynamic matching problems. *American Economic Journal: Microeconomics*, 11(2):55–97.

Kurino, M. (2020). Credibility, efficiency, and stability: A theory of dynamic matching markets. *Japanese Economic Review*, 71:135–165.

Liu, C. (2018). Stability in repeated matching markets. Mimeo.

Liu, Q., Mailath, G. J., Postlewaite, A., and Samuelson, L. (2014). Stable matching with incomplete information. *Econometrica*, 82(2):541–587.

MacDonald, D. E. (2019). Foster care: A dynamic matching approach. Mimeo.

Narita, Y. (2018). Match or mismatch? Learning and inertia in school choice. Mimeo.

Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050.

Pereyra, J. S. (2013). A dynamic school choice model. *Games and Economic Behavior*, 80:100–114.

Roth, A. E. (1984). Stability and polarization of interests in job matching. *Econometrica*, 52(1):47–58.

Roth, A. E. (1985). The college admissions problem is not equivalent to the marriage problem. *Journal of Economic Theory*, 36(2):277–288.

Roth, A. E. (2002). The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378.

Roth, A. E. and Sotomayor, M. A. O. (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, New York.

Samuelson, W. F. and Zeckhauser, R. J. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, 1(1):7–59.

Sasaki, H. and Toda, M. (1996). Two-sided matching problems with externalities. *Journal of Economic Theory*, 70(1):93–108.

Scarf, H. E. (1971). On the existence of a cooperative solution for a general class of N-person games. *Journal of Economic Theory*, 3(2):169–181.

Shapley, L. S. and Shubik, M. (1971). The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130.

Sweeney, S. (2019). Uber, Lyft drivers manipulate fares at Reagan National causing artificial price surges. WJLA News, May 16, 2019. https://wjla.com/news/local/uber-and-lyft-drivers-fares-at-reagan-national (accessed May 20, 2019).

Thakral, N. (2019). Matching with stochastic arrival. *AEA Papers and Proceedings*, 109:209–12.

UCOP [University of California Office of the President] (2019). Fact Sheet: Fall 2019 Freshman and Transfer Admissions. June 26, 2019.

Zhang, M. and Zheng, J. (2016). Multi-period matching with commitment. Mimeo.