Prudence and Stability under Incomplete Information*

Ziwei Wang†

November 24, 2026

Abstract

This paper studies stability in one-to-one matching markets with general contracts and one-sided incomplete information. We accommodate firms' prudence by allowing them to evaluate the status quo and hypothetical deviations under separately chosen beliefs. We define rationalizable α-stability by iterating stability, strong belief in the preceding round, and self-consistency of each firm's conjecture. Self-consistency requires a firm to account for her own inability to block while holding her conjecture fixed. Omitting this requirement can change the outcomes surviving intermediate rounds but leaves the limiting outcome set unchanged. The common limit is the largest self-stabilizing set and is reached after finitely many rounds. We also establish an outcome equivalence characterization: for every fixed information environment with full-support priors, an outcome is rationalizable α-stable if and only if it can be realized by an α-stable matching function in some expansion of that environment. This characterization provides an informational robustness foundation for the solution concept.

Keywords: Matching, contracts, incomplete information, rationalizability, informational robustness.

*This is a companion paper of another paper titled "Rationalizable Stability in Matching with Incomplete Information." Financial support from the National Natural Science Foundation of China (No. 72403188) is acknowledged. The usual disclaimer applies.

†Guanghua School of Management, Peking University. Email: zwang.econ@gmail.com.

1 Introduction

Stability requires an allocation to survive the objections of agents who could benefit by departing from it. Under complete information, the relevant comparisons depend only on payoffs. Under incomplete information, they also depend on beliefs. A firm considering a new match must assess the worker's type, and his willingness to accept a proposed contract can itself be informative. The observed allocation and the absence of other blocking pairs may further restrict which states the firm regards as possible. A notion of incomplete-information stability must therefore specify how these inferences enter the evaluation of deviations.

These considerations raise two distinct consistency questions. First, must a firm evaluate different hypothetical deviations using conditional beliefs derived from one underlying conjecture? This is a question of *Bayesian consistency*. Second, must she maintain the hypothesis that she cannot block under her own conjecture whenever that hypothesis remains compatible with the states under consideration? This is a question of *self-consistency*. The first requirement links beliefs across deviations. The second requires the firm's beliefs to reflect what she can infer from her own inability to block.

In Wang (2026), rationalizable stability imposes both requirements. Each firm's conjecture is a conditional probability system, and the iterative procedure combines stability with strong belief in the outcomes surviving the preceding round and self-consistency. This paper relaxes Bayesian consistency while retaining self-consistency. More precisely, we allow a firm to hold an on-path belief for evaluating the status quo and a separate off-path belief for each hypothetical deviation. Off-path beliefs are indexed by deviations, not only by their inference events: even two deviations that generate the same inference event may be evaluated under different beliefs. This flexibility may reflect firms' prudence, with each deviation assessed under its own pessimistic belief, or their belief that the occurrence of a particular deviation is correlated with unobserved payoff-relevant factors. In either case, beliefs must still respect the information conveyed by the worker's willingness to participate. A related idea appears in *interim perfect equilibrium* (IPE), defined for dynamic games by Penta (2015) and Catonini and Penta (2026). IPE retains Bayesian updating whenever possible but relaxes certain conditional independence restrictions imposed by standard formulations of perfect Bayesian equilibrium (e.g., Fudenberg and Tirole, 1991).[1]

The distinction is useful for understanding the relationship between the iterative approach

[1]The analogy, however, is limited to the treatment of off-path beliefs. First, IPE is an equilibrium notion, in which unexpected moves are defined relative to an equilibrium strategy profile. Our non-equilibrium approach does not fix such a profile, and a pairwise deviation need not be interpreted as an unexpected move at intermediate rounds of reasoning. Second, IPE is associated with backward induction reasoning, under which unexpected moves may be viewed as mistakes, whereas the present framework incorporates a form of forward induction through strong belief in pairwise rationality.

of Liu et al. (2014) and the belief-based equilibrium approach of Liu (2020). Liu et al. (2014) identify a blocking deviation only when the firm gains strictly at every surviving state consistent with her information and the worker's acceptance. Liu (2020) models on-path and off-path beliefs explicitly; his Section IV.E also considers pessimistic off-path beliefs. These approaches suggest a natural question: when beliefs may be chosen separately for different deviations, what does explicit self-consistency add to the set of stable outcomes?

Our first result gives a precise answer. Self-consistency can restrict supporting beliefs and eliminate outcomes earlier, but it does not change the outcomes surviving the entire iteration. We define *rationalizable α-stability* by repeatedly retaining outcomes that are stable under conjectures satisfying strong belief in the preceding round and self-consistency. We then define a simpler procedure that omits self-consistency and prove that the two procedures have the same limit. Thus, self-consistency remains relevant to the interpretation of firms' reasoning, while it can be omitted when computing the limiting outcome set.

This conclusion requires care. Without self-consistency, a conjecture may put probability one on states at which the firm could profitably block under that very conjecture. Such a conjecture fails to incorporate what the firm can infer from her own inability to block. The alternative procedure is nonetheless a well-defined mathematical test. Our result says that each limiting outcome *can* be supported by a conjecture profile satisfying self-consistency, possibly different from the profile supporting it in the simpler procedure. It does not say that all conjectures allowed by the alternative definition already satisfy self-consistency, nor that the two iterations represent the same reasoning at each round. We provide a two-state example that makes this distinction explicit.

The proof uses a fixed-point characterization. A set of outcomes is self-stabilizing if every outcome in the set is stable under conjectures that strongly believe the set. We first show that the simpler procedure converges to the largest self-stabilizing set. The key observation is that beliefs supporting an outcome can be adjusted to strongly believe a larger set without upsetting stability at that outcome. This gives the monotonicity needed to establish the characterization.

We then show that every outcome in a self-stabilizing set survives the iteration of rationalizable α-stability. For each deviation compatible with such a set, we select a deterring belief from a conjecture supporting a state at which the worker accepts. At each round, we retain these deterring beliefs and construct separate off-path beliefs for the remaining deviations so that the resulting conjecture profile supports the outcome and satisfies strong belief and self-consistency. Thus, the largest self-stabilizing set is contained in the limit of rationalizable α-stability, which is itself contained in the limit of the iteration without self-consistency. Since the latter limit is the largest self-stabilizing set, all three coincide by a sandwich argument.

The independence of deviation-specific beliefs is essential to this argument. A belief selected to deter one deviation does not constrain the belief used for another. With Bayesian consistency, as in Wang (2026), these selections may be incompatible. The outcome equivalence therefore reveals a connection between Bayesian consistency and self-consistency: once Bayesian consistency is relaxed in the manner specified here, the additional inference from the firm's own no-blocking can always be accommodated at the limit.

The limiting equivalence also simplifies computation. At a fixed allocation, the iteration without self-consistency can be expressed directly in terms of payoffs at surviving states. A firm can be individually rational under an admissible on-path belief if and only if her status quo payoff is nonnegative at *some* surviving state consistent with her information. Likewise, a deviation accepted by the worker can be deterred if and only if the firm's gain is nonpositive at *some* surviving state consistent with her information and the worker's acceptance. The corresponding beliefs can be chosen as Dirac measures at those states. Thus, each round requires only payoff comparisons, without searching over probability distributions or verifying self-consistency. At each allocation, the iteration terminates within $|\Omega|$ rounds and yields exactly the rationalizable α-stable outcomes.

Our second result connects these iterative predictions to an equilibrium formulation with fixed exogenous beliefs. A matching function specifies an allocation at every state, and agents use the observed allocation to infer which states could have generated it. We fix a subjective prior for each firm and require her on-path belief to be obtained by Bayesian updating after observing the allocation and her information. Off-path beliefs remain separately chosen, subject to their support restrictions. A matching function is α-stable if workers are individually rational, firms are individually rational under their on-path posteriors, and each relevant deviation can be deterred by an admissible off-path belief. For every fixed information environment with full-support priors, we prove that an outcome is rationalizable α-stable if and only if it can be realized by an α-stable matching function in some expansion that preserves the original payoffs and priors. Uninformative expansions already suffice: no additional payoff-relevant information from an expansion is needed.

The construction adds a copy of each state supporting the target allocation. The target allocation is assigned to these copies, while complete-information stable allocations are assigned to the original copies. A small amount of probability is transferred to the added copies according to supporting on-path beliefs, preserving each firm's prior over the original states and ensuring individual rationality under Bayesian updating. Conversely, the projected outcomes of an α-stable matching function on any expansion form a self-stabilizing set, supported by projecting its on-path posteriors and separately selected off-path beliefs. The result gives a precise informational robustness interpretation: outcomes that are not rationalizable cannot be realized in any admissible expansion, while every rationalizable

outcome can be realized in some such expansion. An analyst who knows the original information environment can therefore use rationalizable α-stability to characterize the possible stable outcomes while remaining agnostic about any additional information firms may receive.

Our analysis allows general contracts that may specify transfers, tasks, hours, or other terms of a relationship. Existing pairs may also renegotiate their contracts. The framework thus encompasses standard transferable-utility and non-transferable-utility markets while accommodating a broader range of applications. Following Fisher (2020), we assume a common compact contract set and upper-semicontinuous payoffs, ensuring that a complete-information stable allocation exists at every state. The fixed-point and informational robustness arguments rely on these assumptions only through complete-information existence.

The paper relates to several strands of the literature. Its matching framework includes the classical models of Gale and Shapley (1962) and Shapley and Shubik (1971), with contractual terms treated as part of the relationship as in Hatfield and Milgrom (2005). Its iterative perspective follows rationalizability (Bernheim, 1984; Pearce, 1984), and its strong-belief requirement maintains a hypothesis whenever the relevant inference does not falsify it (Battigalli and Siniscalchi, 2002; Stalnaker, 1998). Related work uses rationalizability to discipline conjectures about future matching opportunities in dynamic markets (Nicolò et al., 2026a) and about the matches formed by other agents following a deviation in markets with externalities (Nicolò et al., 2026b). Pomatto (2022) provides a forward-induction foundation for incomplete-information matching stability through a non-cooperative framework. We work directly with cooperative blocking tests and do not specify a negotiation procedure. Finally, the informational characterization is analogous to results linking rationalizability and equilibrium predictions in non-cooperative games (Brandenburger and Dekel, 1987; Battigalli and Siniscalchi, 2003; Bergemann and Morris, 2017). Our result preserves the given prior beliefs while allowing the information environment to expand, with independent off-path assessments of cooperative deviations.

The rest of the paper is organized as follows. Section 2 presents the model and the criterion of stability. Section 3 defines rationalizable α-stability and the alternative iteration without self-consistency. Section 4 establishes the fixed-point characterization and explains how the alternative procedure simplifies computation. Section 5 gives the informational robustness result. Proofs of the main results appear in the appendix.

2 The Model

2.1 Workers, firms, and contracts

Consider a matching market with finite, nonempty, and disjoint sets of workers I and firms J. Each agent can match with at most one agent on the other side. Let C be a nonempty compact subset of a topological vector space, common to all worker–firm pairs. Each worker i has a type $t_i \in T_i$. If worker i of type t_i matches with firm j under contract $c \in C$, their payoffs are given by the real-valued functions $u_{ij}(t_i, c)$ and $v_{ij}(t_i, c)$, respectively. For every pair (i, j) and every type $t_i \in T_i$, both payoff functions are upper-semicontinuous in c. Unmatched agents receive zero. Contracts may specify transfers, tasks, hours, or other terms of a relationship.

A matching $\mu : I \cup J \to I \cup J$ satisfies $\mu(i) \in J \cup \{i\}$ for every worker i, $\mu(j) \in I \cup \{j\}$ for every firm j, and $\mu(\mu(k)) = k$ for every agent $k \in I \cup J$. An agent k is unmatched if $\mu(k) = k$. An allocation $x = (\mu, \mathbf{c})$ specifies a matching and a contract $c_{ij} \in C$ for every pair (i, j) with $\mu(i) = j$. Write X for the set of allocations.

There is a finite, nonempty set of underlying states Ω. A mapping $\tau : \Omega \to T$, where $T = \times_{i \in I} T_i$, assigns a worker type profile to each state. Write $\tau_i(\omega) = (\tau(\omega))_i$. Workers know their own types and hence their own payoffs. Firms may be uncertain about worker types, including the type of their current employee. An *outcome* is a pair $(\omega, x) \in \mathcal{O} := \Omega \times X$. With a slight abuse of notation, we write $u_{ij}(\omega, c) = u_{ij}(\tau_i(\omega), c)$ and $v_{ij}(\omega, c) = v_{ij}(\tau_i(\omega), c)$. Throughout, the status quo payoffs $u_{i\mu(i)}(\omega, c_{i\mu(i)})$ and $v_{\mu(j)j}(\omega, c_{\mu(j)j})$ are understood to be zero when worker i or firm j, respectively, is unmatched.

For each pair (i, j) and allocation x, let $C_{ij}(x) \subseteq C$ be the set of contracts available to that pair as deviations. The unrestricted case is $C_{ij}(x) = C$. Feasibility may depend on the publicly observed allocation, but not on the unobserved state. A feasible pairwise deviation (or a deviation for brevity) is a triple $(i, j, \hat{c})$ with $\hat{c} \in C_{ij}(x)$. We allow an existing pair to change its contract: with general contracts, such a change can strictly benefit both parties.

This framework includes two familiar cases. In a transferable-utility market, let $C = [-\bar{p}, \bar{p}]$ for some $\bar{p} > 0$ and set

$$u_{ij}(t_i, c) = a_{ij}(t_i) + c, \qquad v_{ij}(t_i, c) = b_{ij}(t_i) - c.$$

The interval can be chosen sufficiently large to retain every individually rational outcome and every deviation that could block such an outcome.[2] In a market without transfers or other

[2]Let $P > 0$ bound the absolute values of $a_{ij}(\tau_i(\omega))$ and $b_{ij}(\tau_i(\omega))$ over all i, j, and ω. Workers' individual rationality and firms' expected individual rationality imply that every status quo transfer lies in $[-P, P]$. Status quo payoffs are therefore at least $-2P$ at every state. A transfer below $-3P$ cannot attract the worker at any state, and one above $3P$ cannot benefit the firm at any state. Thus, $\bar{p} \geq 3P$ suffices.

contractual choices, C can be a singleton. The contract set may be infinite; the finite-round bound established below relies on the finiteness of Ω.

An outcome (ω, x) with $x = (\mu, \mathbf{c})$ is *complete-information stable* if all agents receive nonnegative payoffs and there is no feasible deviation $(i, j, \hat{c})$ such that

$$u_{ij}(\omega, \hat{c}) > u_{i\mu(i)}(\omega, c_{i\mu(i)}) \quad \text{and} \quad v_{ij}(\omega, \hat{c}) > v_{\mu(j)j}(\omega, c_{\mu(j)j}).$$

Under the maintained assumptions on the contract space and payoffs, complete-information existence follows from Fisher (2020).

Lemma 1 (Existence; Fisher, 2020)**.** *At every state $\omega \in \Omega$, there exists a complete-information stable allocation.*

Proof. Fix $\omega \in \Omega$. By Fisher (2020), a complete-information stable allocation exists when every contract in C is available to every pair. Restricting deviations to $C_{ij}(x) \subseteq C$ cannot introduce a blocking pair, so the same allocation is stable under our feasibility restrictions. □

We use compactness and upper semicontinuity only to obtain this existence result.

2.2 Information and conjectures

We assume that the allocation is publicly observed. Firm j's information at an outcome is described by an *information correspondence* $\Pi_j : \mathcal{O} \rightrightarrows \Omega$ satisfying the following conditions for every $(\omega, x) \in \mathcal{O}$:

(i) $\omega \in \Pi_j(\omega, x)$;

(ii) if $\omega' \in \Pi_j(\omega, x)$, then $\Pi_j(\omega', x) = \Pi_j(\omega, x)$.

Thus, for each fixed allocation, the information sets form a partition of Ω. The partition may vary with the allocation, including its contractual terms. Write $\mathcal{I} = \langle \Omega, \tau, (\Pi_j)_{j \in J} \rangle$ for the *information structure*. The market primitives and the information structure are common knowledge.

When evaluating a hypothetical deviation $(i, j, \hat{c})$, firm j takes into account worker i's willingness to participate. The corresponding inference event is

$$D_j(\omega, x; i, \hat{c}) = \{\omega' \in \Pi_j(\omega, x) : u_{ij}(\omega', \hat{c}) > u_{i\mu(i)}(\omega', c_{i\mu(i)})\}.$$

If this set is empty, the worker rejects the deviation at every state consistent with the firm's information. Such a deviation cannot block and requires no belief specification.

For every nonempty $E \subseteq \Omega$, let $\Delta(E)$ denote the set of probability distributions supported on E.

Definition 1 (Conjectures)**.** A *conjecture* ν_j of firm j at (ω, x) consists of an on-path belief $\nu_j^0 \in \Delta(\Pi_j(\omega, x))$ and an off-path belief $\nu_j^{(i,\hat{c})} \in \Delta(D_j(\omega, x; i, \hat{c}))$ for every feasible deviation $(i, j, \hat{c})$ with $D_j(\omega, x; i, \hat{c}) \neq \varnothing$. A *conjecture profile* is $\nu = (\nu_j)_{j \in J}$.

The on-path belief is used to evaluate the status quo, whereas each off-path belief is used to evaluate a particular deviation conditional on the worker's willingness to participate. These beliefs may be specified separately. In particular, two deviations may generate the same inference event and nevertheless be evaluated under different off-path beliefs. Moreover, $\nu_j^{(i,\hat{c})}$ need not be obtained by conditioning ν_j^0 on $D_j(\omega, x; i, \hat{c})$, even when this event has positive probability under ν_j^0.

This freedom accommodates a form of *prudence*. A firm may assess each proposed deviation pessimistically, using a conjecture specific to that deviation. An analyst who does not wish to specify how firms interpret hypothetical deviations can therefore leave these beliefs unrestricted, subject to their support restrictions. This gives a permissive test of stability: an outcome is excluded only if no admissible conjecture profile can support it, so its exclusion does not depend on a particular choice of off-path beliefs.

For a distribution π on a subset of Ω, we write $\pi[E]$ for the probability assigned to $E \subseteq \Omega$ and $\mathbb{E}_\pi[f]$ for the expectation of f. All beliefs are supported on finite sets. We do not impose an exogenous prior or a common-prior restriction.

2.3 Blocking and stability

Definition 2 (Blocking pairs)**.** A feasible pairwise deviation $(i, j, \hat{c})$ *blocks* (ω, x) under ν_j if

$$u_{ij}(\omega, \hat{c}) > u_{i\mu(i)}(\omega, c_{i\mu(i)}), \tag{1}$$

$$\mathbb{E}_{\nu_j^{(i,\hat{c})}}[v_{ij}(\cdot, \hat{c})] > \mathbb{E}_{\nu_j^{(i,\hat{c})}}[v_{\mu(j)j}(\cdot, c_{\mu(j)j})]. \tag{2}$$

The first inequality requires the worker to strictly benefit at the true state. The second requires the firm to expect a strict gain, evaluating the proposed contract and the status quo under the same off-path belief.

Definition 3 (Stability)**.** An outcome (ω, x) is *stable with respect to* a conjecture profile ν if

(i) for every worker i, $u_{i\mu(i)}(\omega, c_{i\mu(i)}) \geq 0$;

(ii) for every firm j, $\mathbb{E}_{\nu_j^0}[v_{\mu(j)j}(\cdot, c_{\mu(j)j})] \geq 0$;

(ii) no feasible pairwise deviation $(i, j, \hat{c})$ blocks (ω, x) under ν_j.

Conditions (i) and (ii) are individual rationality. Condition (iii) is no-blocking. A firm's on-path belief governs her willingness to maintain the existing match, whereas her off-path belief governs the comparison between that match and a proposed deviation.

To express what firm j can infer from her own inability to block, define

$$NB_j(\omega, x; \nu_j) = \{\omega' \in \Pi_j(\omega, x) : \text{no feasible deviation } (i, j, \hat{c}) \text{ blocks } (\omega', x) \text{ under } \nu_j\}.$$

The entire conjecture ν_j is held fixed in this definition. This is well-defined because, at a fixed allocation, states in the same information cell have the same feasible deviations and inference events.

A useful consequence is that if a deviation $(i, j, \hat{c})$ blocks (ω, x) under ν_j, then it also blocks (ω', x) under the same conjecture for every $\omega' \in D_j(\omega, x; i, \hat{c})$. The firm's expected gain is unchanged, and the worker strictly benefits at every state in this event.

3 Solution Concept

3.1 Strong belief, self-consistency, and rationalizable α-stability

Stability under some conjecture profile is a preliminary requirement. Firms also use the stability of the observed allocation to reason about which states remain possible. For a set of outcomes $R \subseteq \mathcal{O}$, write

$$R_x = \{\omega \in \Omega : (\omega, x) \in R\}.$$

We first formalize the requirement that firm j maintain her belief in R_x whenever the information she uses does not rule it out.

Definition 4 (Strong belief)**.** A conjecture ν_j at (ω, x) *strongly believes* R_x if

(i) whenever $R_x \cap \Pi_j(\omega, x) \neq \varnothing$, we have $\nu_j^0[R_x] = 1$;

(ii) for every deviation $(i, j, \hat{c})$ such that $R_x \cap D_j(\omega, x; i, \hat{c}) \neq \varnothing$, we have $\nu_j^{(i,\hat{c})}[R_x] = 1$.

The second condition requires the firm to retain the hypothesis R_x after learning that a worker accepts a proposed contract, provided that acceptance is compatible with R_x. If the inference event is disjoint from R_x, the hypothesis is falsified and strong belief imposes no additional restriction on the off-path belief. This is the strong-belief idea of Battigalli and Siniscalchi (2002) and Stalnaker (1998), applied here to separately specified beliefs rather than a conditional probability system.

A further restriction arises from the firm's own inability to block. The firm knows the conjecture she uses. Thus, when inferring which states remain possible, she should assess her own no-blocking while holding this conjecture fixed, rather than using a different conjecture that might support the same allocation (see Wang, 2026, for a detailed discussion).

Definition 5 (Self-consistency)**.** A conjecture ν_j at (ω, x) is *self-consistent relative to* R_x if

(i) whenever $R_x \cap NB_j(\omega, x; \nu_j) \neq \varnothing$, we have $\nu_j^0[NB_j(\omega, x; \nu_j)] = 1$;

(ii) for every deviation $(i, j, \hat{c})$ such that $R_x \cap NB_j(\omega, x; \nu_j) \cap D_j(\omega, x; i, \hat{c}) \neq \varnothing$, we have $\nu_j^{(i,\hat{c})}[NB_j(\omega, x; \nu_j)] = 1$.

When strong belief and self-consistency both apply to an off-path belief, that belief assigns probability one to $R_x \cap NB_j(\omega, x; \nu_j)$. Thus, the firm concentrates on states that remain possible under R and at which she herself cannot block under the same conjecture. In particular, if $(\omega, x) \in R$ is stable under ν, then $\omega \in R_x \cap NB_j(\omega, x; \nu_j)$ for every firm j, so the on-path requirement of self-consistency applies.

Definition 6 (Rationalizable α-stability)**.** Let $R^0 = \mathcal{O}$. For each $n \geq 1$, define R^n as the set of outcomes $(\omega, x) \in R^{n-1}$ for which there exists a conjecture profile ν such that

(i) (ω, x) is stable with respect to ν;

(ii) every ν_j strongly believes R_x^{n-1};

(iii) every ν_j is self-consistent relative to R_x^{n-1}.

Define $R^\infty = \bigcap_{n \geq 0} R^n$. We call its elements *rationalizable α-stable outcomes*.

The iteration removes outcomes that cannot be supported once firms account for the preceding round and their own no-blocking. This is a non-equilibrium approach: it does not presuppose a commonly known rule specifying the allocation at every underlying state. Firms use stability to restrict which states could support the observed allocation, and retain these hypotheses after a hypothetical deviation whenever they remain compatible with the worker's willingness to participate. This iterative use of strong belief parallels the reasoning behind strong rationalizability in dynamic games (Battigalli and Siniscalchi, 2002). The related notion of rationalizable stability in Wang (2026) additionally imposes Bayesian consistency: each firm's on-path and off-path beliefs must belong to a single conditional probability system.

The qualifier α reflects the prudent evaluation of potential deviations. To see the idea, fix a nonempty set of states compatible with a deviation and allow the firm to hold any belief supported on this set. The deviation yields a strict expected gain under every such belief if and only if it strictly benefits the firm at every state in the set. Otherwise, a Dirac belief at a state where the firm is no better off deters the deviation. In our iterative procedure, however, the beliefs used to deter deviations must also satisfy the strong-belief and self-consistency requirements of each round. Section 4.2 makes the resulting state-by-state test explicit. Kotowski (2024) applies a related robustness logic to plausible market continuations in multiperiod matching; here, uncertainty concerns the underlying state. Related approaches use rationalizability to discipline conjectures about the consequences of deviations in dynamic matching (Nicolò et al., 2026a) and matching with externalities (Nicolò et al., 2026b).

3.2 An alternative iteration without self-consistency

Consider the following simpler procedure. Let $\widehat{R}^0 = \mathcal{O}$. For each $n \geq 1$, define $\widehat{R}^n$ as the set of outcomes $(\omega, x) \in \widehat{R}^{n-1}$ for which there exists a conjecture profile ν at (ω, x) such that

(i) (ω, x) is stable with respect to ν;

(ii) every ν_j strongly believes $\widehat{R}^{n-1}_x$.

Define $\widehat{R}^\infty = \bigcap_{n\geq 0} \widehat{R}^n$. This procedure differs from Definition 6 only by omitting self-consistency.

The omission matters for how firms reason. Strong belief requires a firm to concentrate on outcomes surviving the preceding round whenever they remain compatible with the information she uses. It does not require her own no-blocking to hold at those states under her current conjecture. A firm may therefore assign probability one to a state at which she would herself block under that same conjecture. Such a belief fails to incorporate what she can infer from her own inability to block. The simpler procedure is similar in spirit to the iterative approach of Liu et al. (2014). The following example shows that omitting self-consistency can change the outcomes retained even in the first round.

Example 1. There is one worker i, one firm j, and two states $\Omega = \{\ell, h\}$. Contracts are transfers $p \in C = [-P, P]$, where $P \geq 30$, and every contract is available to the pair. The worker's and firm's matching values before transfers are

	ℓ	h
a	-1	-4
b	3	10

Let x^0 leave both agents unmatched, and suppose the firm has no information at this allocation. Both status quo payoffs are zero, so individual rationality holds under every conjecture. The inference events are

$$D_j(\omega, x^0; i, p) = \begin{cases} \varnothing, & p \leq 1, \\ \{\ell\}, & 1 < p \leq 4, \\ \{\ell, h\}, & p > 4. \end{cases}$$

The iteration without self-consistency. At state ℓ, payment $p = 2$ strictly benefits both parties. Its inference event is $\{\ell\}$, so every admissible off-path belief assigns probability one to ℓ. Hence, $(\ell, x^0) \notin \widehat{R}^1$. At state h, the worker accepts only payments $p > 4$. Let the firm assign probability one to ℓ both on path and after every deviation with a nonempty inference event. At every payment the worker accepts, the firm's expected payoff is $3 - p < 0$. This conjecture supports (h, x^0) in $\widehat{R}^1$, so $\widehat{R}^1_{x^0} = \{h\}$. In the second round, strong belief requires

the off-path belief at $p = 5$ to assign probability one to h. Both agents then strictly benefit from the deviation, so $\widehat{R}^2_{x^0} = \varnothing$.

The iteration with self-consistency. State ℓ is again eliminated immediately. Suppose a conjecture ν_j supports (h, x^0) in R^1. Since payment $p = 2$ blocks at ℓ under every conjecture, whereas no deviation blocks at h under ν_j, we must have $NB_j(h, x^0; \nu_j) = \{h\}$. Because $R^0_{x^0} = \Omega$ and the inference event at $p = 5$ contains h, self-consistency requires $\nu_j^{(i,5)}[h] = 1$. The deviation then blocks at h, a contradiction. Thus, $R^1_{x^0} = \varnothing$. ♢

Intuitively, the conjecture supporting (h, x^0) in the first round of the simpler procedure allows the firm to reason as follows:

> "I do not want to hire the worker at a payment above four because I believe that his type is ℓ. But if his type were ℓ, we would both gain from matching at a payment of two. My belief therefore cannot explain why no payment allows us to form a blocking pair. I nevertheless use this very belief to justify the stability of the unmatched allocation."

The firm thus assigns probability one to a state incompatible with her own no-blocking under the *same conjecture*. The simpler iteration is mathematically well-defined, but its supporting beliefs need not incorporate this implication of no-blocking. In this example, the second round removes the outcome that self-consistency rules out immediately.

Our first result shows that this coincidence of limiting predictions is general: $R^\infty = \widehat{R}^\infty$. The statement is about outcomes, not about the conjectures used to support them. In particular, it does not say that an arbitrary conjecture supporting an outcome in the simpler procedure is self-consistent. The proof constructs suitable replacement conjectures. This distinction explains why self-consistency can be essential to the interpretation of the reasoning while being dispensable for computing the limiting outcome set.

4 A Fixed-Point Characterization

4.1 Self-stabilizing sets

A set of outcomes is self-stabilizing if every outcome in the set can be supported by conjectures that strongly believe the set. Such a set describes outcomes that can be justified together, without referring to a particular round of the iteration.

Definition 7 (Self-stabilizing sets)**.** A nonempty set $F \subseteq \mathcal{O}$ is *self-stabilizing* if, for every $(\omega, x) \in F$, there exists a conjecture profile ν at (ω, x) such that (ω, x) is stable with respect to ν and every ν_j strongly believes F_x.

This definition does not impose self-consistency. The following elementary properties parallel those in Liu et al. (2014).

Lemma 2.

(i) A singleton $\{(\omega, x)\}$ is self-stabilizing if and only if (ω, x) is complete-information stable.

(ii) The union of any nonempty collection of self-stabilizing sets is self-stabilizing.

We next show that the supporting conjectures can be chosen so that each firm cannot block at *any* state in the set consistent with her information.

Lemma 3. *Let F be self-stabilizing and fix $(\omega, x) \in F$. There exists a conjecture profile ν supporting (ω, x) in F under Definition 7 such that, for every firm j,*

$$F_x \cap \Pi_j(\omega, x) \subseteq NB_j(\omega, x; \nu_j).$$

In particular, every ν_j is self-consistent relative to F_x.

The idea is simple. Consider a deviation whose inference event intersects F_x. Select a state in this intersection and a conjecture profile supporting that state in F. Since the worker accepts the deviation there, the supporting off-path belief must deter the firm. The selected state and the original state lie in the same information cell, so the inference event is unchanged. We can therefore use the selected belief to evaluate this deviation at the original outcome. Repeating this selection for each such deviation rules out the firm's blocking at every state in $F_x \cap \Pi_j(\omega, x)$, while leaving the on-path belief unchanged. Crucially, the off-path belief selected for one deviation places no restriction on the belief selected for another.[3]

The lemma gives more than self-consistency relative to F_x: it provides a *single* conjecture under which the firm cannot block at any state in $F_x \cap \Pi_j(\omega, x)$. Nevertheless, it does not immediately establish that F survives the iteration in Definition 6. At an intermediate round, firms must strongly believe the potentially larger set R_x^{n-1} and satisfy self-consistency relative to that set. The proof below addresses these additional requirements by adjusting beliefs at deviations whose inference events are disjoint from F_x.

Proposition 1 (Fixed point, outcome equivalence, and existence)**.**

(i) $R^n \subseteq \widehat{R}^n$ for every $n \geq 0$ and $R^\infty = \widehat{R}^\infty$.

(ii) R^∞ is self-stabilizing and contains every self-stabilizing set.

(iii) Both procedures terminate after at most $|\Omega|$ rounds: $R^{|\Omega|} = \widehat{R}^{|\Omega|} = R^\infty$.

(iv) Every complete-information stable outcome belongs to R^∞. In particular, for every $\omega \in \Omega$, there exists $x \in X$ such that $(\omega, x) \in R^\infty$.

[3]Such a selection need not be feasible if Bayesian consistency is imposed across beliefs; see Wang (2026).

Part (i) shows that self-consistency may eliminate outcomes earlier without changing the limiting predictions. Part (ii) gives a characterization without reference to the iteration: an outcome survives precisely when it belongs to a self-stabilizing set. Part (iii) follows because, for each fixed allocation, either procedure updates a decreasing subset of Ω. Once a round leaves this subset unchanged, it remains unchanged in all subsequent rounds. Thus, at most $|\Omega|$ rounds are needed, even when the sets of contracts and allocations are infinite. Part (iv) establishes existence using Lemma 1.

The proof proceeds in two steps. First, we show that the simpler procedure is monotone in the preceding set. Enlarging that set may require adjusting beliefs at deviations that become compatible with it. These deviations are rejected by the worker at the outcome being supported, so the adjustments preserve stability. The limit of the simpler procedure is therefore the largest self-stabilizing set. Second, take any self-stabilizing set $F \subseteq R^{n-1}$. Lemma 3 provides beliefs that deter every deviation compatible with F_x. Preserve those beliefs, adjust the remaining beliefs to ensure strong belief in R_x^{n-1}, and then refine them to obtain self-consistency relative to that set. Previously blocking deviations remain blocking throughout this refinement, so the relevant no-blocking sets decrease and the refinement stops after finitely many steps. This proves $F \subseteq R^n$ and, by induction, establishes the converse inclusion at the limit. Appendix A.1 gives the constructions and verifies each condition.

4.2 Computing R^∞

Proposition 1 allows us to compute all rationalizable α-stable outcomes using the iteration without self-consistency. This procedure admits a simple state-by-state formulation. Fix a set $R \subseteq \mathcal{O}$ and an outcome $(\omega, x) \in R$. There exists a conjecture profile ν such that (ω, x) is stable with respect to ν and every ν_j strongly believes R_x if and only if the following conditions hold:

(i) For every worker i, $u_{i\mu(i)}(\omega, c_{i\mu(i)}) \geq 0$;

(ii) For every firm j, there exists $\omega' \in R_x \cap \Pi_j(\omega, x)$ such that $v_{\mu(j)j}(\omega', c_{\mu(j)j}) \geq 0$;

(iii) For every deviation $(i, j, \hat{c})$ accepted by the worker at ω, there exists $\omega' \in R_x \cap D_j(\omega, x; i, \hat{c})$ such that $v_{ij}(\omega', \hat{c}) \leq v_{\mu(j)j}(\omega', c_{\mu(j)j})$.

The sets in conditions (ii) and (iii) are nonempty because they contain ω. Necessity follows from individual rationality and no-blocking: each required inequality in expectation must hold at some state in the support of the relevant belief. For sufficiency, choose each firm's on-path belief to be a Dirac measure at a state satisfying (ii). For each deviation accepted at ω, choose the off-path belief to be a Dirac measure at a state satisfying (iii). For every remaining deviation with a nonempty inference event, choose any belief supported on $R_x \cap D_j(\omega, x; i, \hat{c})$

if this intersection is nonempty, and any belief supported on $D_j(\omega, x; i, \hat{c})$ otherwise. These choices satisfy strong belief and make (ω, x) stable.

Thus, for a fixed allocation x, start with all states in Ω and repeatedly retain those satisfying (i)–(iii). By Proposition 1, this procedure reaches R_x^∞ after at most $|\Omega|$ rounds. In each round, the analyst need not solve for self-consistent conjectures or impose the no-blocking restriction within each candidate belief. This simplifies the computation of the limiting outcome set. If self-consistent supporting conjectures are also needed, they can be constructed at the limit using Lemma 3.

4.3 Relation to Liu et al. (2014)

The test in condition (iii) above makes the connection to Liu et al. (2014) transparent. For a deviation accepted by the worker at the true state, their definition identifies a block if and only if the firm gains strictly at *every* surviving state consistent with her information and the worker's willingness to deviate. Thus, an outcome admits no blocking pair in their sense if and only if every such deviation has a surviving state at which the firm is no better off. This is precisely condition (iii).

Under their transferable-utility primitives and information structure, each firm observes her current employee's type, so her status quo payoff is constant on her information cell. Condition (ii) therefore reduces to firms' ex post individual rationality. Under these primitives, the process we describe in Section 4.2 reduces to the one in Liu et al. (2014). Our framework nonetheless allows general contracts and arbitrary partitional information, including uncertainty about the current employee. Our analysis further shows that every limiting outcome admits self-consistent supporting conjectures, even though self-consistency need not be imposed in the procedure used to compute those outcomes.

This equivalence between the limiting predictions with and without self-consistency relies on more than the absence of a common prior. It uses the ability to choose off-path beliefs separately even when two deviations generate the same inference event. In the Bayesian formulation of Wang (2026), each firm's conditional beliefs must belong to a single conditional probability system. A belief chosen to deter one deviation can then constrain the evaluation of another, and the replacement argument in Lemma 3 need not be feasible. The redundancy of self-consistency established here therefore depends on the freedom to specify beliefs separately across deviations; it does not extend automatically to Bayesian-consistent conjectures.

5 Informational Robustness

The non-equilibrium approach asks which outcomes can be justified without specifying how the allocation is selected at every state. An equilibrium approach instead starts with a matching function that describes this selection and is correctly understood by agents. We

now connect these two approaches through an outcome equivalence result, analogous to those linking rationalizability and equilibrium predictions in noncooperative games (Brandenburger and Dekel, 1987; Battigalli and Siniscalchi, 2003; Bergemann and Morris, 2017). For the equilibrium analysis, we fix firms' exogenous beliefs and require them to update these beliefs after observing the allocation. Beliefs following hypothetical deviations remain separately chosen.

5.1 Information environments

An *information environment* is a pair $\mathcal{E} = \langle \mathcal{I}, \beta^0 \rangle$, where $\mathcal{I}$ is the information structure and $\beta^0 = (\beta_j^0)_{j \in J}$ is a profile of exogenous priors $\beta_j^0 \in \Delta(\Omega)$. These priors are fixed primitives of the equilibrium analysis and need not coincide across firms. For every $E, B \subseteq \Omega$ with $\beta_j^0[B] > 0$, the conditional belief is given by Bayes' rule:

$$\beta_j^0[E \mid B] = \frac{\beta_j^0[E \cap B]}{\beta_j^0[B]}.$$

All conditioning events relevant to our characterization result have positive probability, so these conditional beliefs are well defined. Note that the iterative definition of R^∞ does not depend on β^0.

5.2 Stable matching functions

A *matching function* is a map $M : \Omega \to X$. Agents know this function and observe the realized allocation. If $M(\omega) = x$, the states consistent with that observation form the set

$$M^{-1}(x) = \{\omega' \in \Omega : M(\omega') = x\}.$$

Firm j also uses her information $\Pi_j(\omega, x)$. Her on-path belief is therefore $\beta_j^0[\cdot \mid M^{-1}(x) \cap \Pi_j(\omega, x)]$, provided that the conditioning event has positive probability. We require this at every state, so all on-path beliefs are obtained by ordinary Bayesian updating.

When evaluating a hypothetical deviation $(i, j, \hat{c})$, firm j also takes into account worker i's willingness to participate. The states consistent with the allocation, the firm's information, and the worker's acceptance form $M^{-1}(x) \cap D_j(\omega, x; i, \hat{c})$. Whenever this set is nonempty, the firm may choose any belief supported on it, even if the set has zero prior probability. This belief need not be obtained by updating β_j^0 or her on-path belief and may differ across deviations even when they generate the same inference event. As before, this freedom in specifying off-path beliefs reflects the firm's prudence in evaluating hypothetical deviations.

Definition 8 (α-Stable matching functions)**.** A matching function M is *α-stable* on $\mathcal{E}$ if, for every $\omega \in \Omega$ and $x = M(\omega) = (\mu, \mathbf{c})$, the following conditions hold:

(i) For every worker i, we have $u_{i\mu(i)}(\omega, c_{i\mu(i)}) \geq 0$;

(ii) For every firm j, we have $\beta_j^0[M^{-1}(x) \cap \Pi_j(\omega, x)] > 0$ and $\mathbb{E}_{\beta_j^0}[v_{\mu(j)j}(\cdot, c_{\mu(j)j}) \mid M^{-1}(x) \cap \Pi_j(\omega, x)] \geq 0$;

(iii) For every deviation $(i, j, \hat{c})$ such that $M^{-1}(x) \cap D_j(\omega, x; i, \hat{c}) \neq \varnothing$, there exists $\beta_j^{(i,\hat{c})} \in \Delta(M^{-1}(x) \cap D_j(\omega, x; i, \hat{c}))$ such that $\mathbb{E}_{\beta_j^{(i,\hat{c})}}[v_{ij}(\cdot, \hat{c})] \leq \mathbb{E}_{\beta_j^{(i,\hat{c})}}[v_{\mu(j)j}(\cdot, c_{\mu(j)j})]$.

Condition (ii) requires a well-defined on-path posterior under which the firm is individually rational. Condition (iii) requires that each relevant deviation can be deterred by a belief under which the firm does not expect to gain. The dependence of $\beta_j^{(i,\hat{c})}$ on the observed allocation and the firm's information cell is left implicit.

Condition (iii) also covers deviations that the worker might reject at the actual state ω, since $D_j(\omega, x; i, \hat{c})$ need not contain ω. This adds no restriction beyond requiring no-blocking at every state. Indeed, take any $\omega' \in M^{-1}(x) \cap D_j(\omega, x; i, \hat{c})$. The worker accepts the deviation at ω', the allocation is still x, and $\Pi_j(\omega', x) = \Pi_j(\omega, x)$. The inference event and the firm's payoff comparison are therefore unchanged. No-blocking at ω' thus requires the existence of a belief satisfying condition (iii).

This formulation is related to the pessimistic beliefs in Liu (2020) (Section IV.E): on-path beliefs are updated from exogenous beliefs, while hypothetical deviations can be evaluated pessimistically. Here, firms may have different exogenous priors. Moreover, Liu's blocking criterion requires the firm to prefer the deviation to both the status quo match and becoming unmatched. Our comparison is only with the status quo match, as in Definition 2.

5.3 Expansions and an Informational Characterization

An analyst may know the baseline information environment $\mathcal{E}$ without knowing whether firms receive additional information. We represent these possibilities by expansions that preserve the underlying payoffs and exogenous beliefs.

Definition 9 (Expansions)**.** An *expansion* of $\mathcal{E} = \langle \mathcal{I}, \beta^0 \rangle$ with projection ρ consists of an information environment $\bar{\mathcal{E}} = \langle \bar{\mathcal{I}}, \bar{\beta}^0 \rangle$, where $\bar{\mathcal{I}} = \langle \bar{\Omega}, \bar{\tau}, (\bar{\Pi}_j)_{j \in J} \rangle$ has a finite state space, and a surjective map $\rho : \bar{\Omega} \to \Omega$ satisfying the following conditions:

(i) For every $\bar{\omega} \in \bar{\Omega}$, $\bar{\tau}(\bar{\omega}) = \tau(\rho(\bar{\omega}))$;

(ii) For every $\bar{\omega} \in \bar{\Omega}$, allocation x, and firm j, $\bar{\Pi}_j(\bar{\omega}, x) \subseteq \rho^{-1}(\Pi_j(\rho(\bar{\omega}), x))$;

(iii) For every firm j and $E \subseteq \Omega$, $\bar{\beta}_j^0[\rho^{-1}(E)] = \beta_j^0[E]$.

The expansion is *uninformative* if the inclusion in (ii) is an equality for every $\bar{\omega}$, x, and j.

The projection ρ identifies the original state underlying an expanded state. Condition (i) preserves payoffs, and condition (ii) permits additional information while retaining the firm's original information. Condition (iii) preserves each firm's prior over original states. By

Bayes' rule, it also preserves conditional beliefs when both events concern only original states and the conditioning event has positive probability.

In an uninformative expansion, at any fixed allocation, a firm cannot distinguish expanded states whose projections belong to the same original information cell. The expanded state space may nevertheless contain several copies of an original state, and a matching function may assign different allocations to these copies. Observing the allocation can therefore convey information even though the expansion adds no information through the firms' information correspondences.

Proposition 2 (Informational robustness)**.** *Fix an information environment $\mathcal{E} = \langle \mathcal{I}, \beta^0 \rangle$ in which every β_j^0 has full support.*[4] *For every outcome $(\omega, x) \in \mathcal{O}$, the following statements are equivalent:*

(i) $(\omega, x) \in R^\infty$.

(ii) There exist an uninformative expansion $\bar{\mathcal{E}}$ of $\mathcal{E}$ with projection ρ, an α-stable matching function M on $\bar{\mathcal{E}}$, and a state $\bar{\omega} \in \rho^{-1}(\omega)$ such that $M(\bar{\omega}) = x$.

(iii) There exist an expansion $\bar{\mathcal{E}}$ of $\mathcal{E}$ with projection ρ, an α-stable matching function M on $\bar{\mathcal{E}}$, and a state $\bar{\omega} \in \rho^{-1}(\omega)$ such that $M(\bar{\omega}) = x$.

The proposition gives a precise *informational robustness* interpretation of R^∞. Suppose an analyst knows $\mathcal{E}$ but does not know whether firms receive additional information. An outcome outside R^∞ cannot be generated by an α-stable matching function in any expansion consistent with this environment. Conversely, every outcome in R^∞ can be generated in some such expansion. Thus, R^∞ describes exactly the outcomes compatible with α-stability when the analyst leaves the expansion unspecified.

Part (ii) strengthens the converse: uninformative expansions already suffice to generate every outcome in R^∞. The realization of a rationalizable outcome does not require firms to receive additional information through their information correspondences, although they may still learn from the observed allocation. The characterization holds for every full-support profile of original priors β^0, even though the iterative definition of R^∞ does not refer to these beliefs. The supporting expansion and matching function may depend on both β^0 and the outcome to be realized.

The proof uses the self-stabilizing property of R^∞. To realize $(\omega, x) \in R^\infty$, add a copy of each state in R_x^∞ and assign x to these copies. Assign complete-information stable allocations

[4] Full support allows us to work with ordinary probability distributions. An analogous characterization can be obtained without this assumption by using conditional probability systems and allowing on-path beliefs to be specified following zero-probability events. This extension requires a more involved construction but leaves the economic interpretation unchanged.

to the original copies. A small amount of probability is transferred to the added copies according to supporting on-path beliefs, preserving each original prior and ensuring individual rationality after the allocation is observed. Conversely, projecting the outcomes of an α-stable matching function on any expansion gives a self-stabilizing set. Proposition 1 then places every projected outcome in R^∞. Appendix A.2 provides the details.

6 Conclusion

We propose a non-equilibrium notion of stability that incorporates firms' prudence in evaluating hypothetical deviations. Each firm may assess different deviations under different off-path beliefs, even when they generate the same inference event. In the terminology of Liu (2020), we relax Bayesian consistency across on-path and off-path beliefs. Self-consistency addresses a different aspect of reasoning: it requires a firm to account for her own inability to block while holding her conjecture fixed. We show that, when beliefs can be chosen separately for different deviations, self-consistency can be omitted from the iterative procedure without changing its limiting outcomes. Thus, a simpler iteration suffices for computation, while every surviving outcome still admits self-consistent supporting conjectures.

The common limit also coincides with the outcomes generated by α-stable matching functions across expansions of any fixed information environment with full-support priors. Each expansion preserves the original priors, and on-path beliefs are obtained by Bayesian updating at events of positive probability. Off-path beliefs, nevertheless, remain separately chosen. Uninformative expansions already suffice for the characterization. This result provides a foundation for using rationalizable α-stability to make informationally robust predictions when the analyst knows the original environment but is agnostic about the additional information firms may receive.

A Appendix

A.1 Proofs for Section 4

We begin with a monotonicity observation used below. Let $R \subseteq R' \subseteq \mathcal{O}$ and $(\omega, x) \in R$. Suppose a conjecture profile ν supports stability of (ω, x) and strongly believes R_x. We can construct a profile ν' that preserves stability and strongly believes R'_x. Keep all on-path beliefs unchanged. For each deviation $(i, j, \hat{c})$ such that $R_x \cap D_j(\omega, x; i, \hat{c}) = \varnothing$ but $R'_x \cap D_j(\omega, x; i, \hat{c}) \neq \varnothing$, set

$$(\nu'_j)^{(i,\hat{c})} = \text{Unif}(R'_x \cap D_j(\omega, x; i, \hat{c})).$$

Keep all remaining off-path beliefs unchanged. The resulting profile strongly believes R'_x and preserves individual rationality. Moreover, every deviation accepted at ω has an inference event containing $\omega \in R_x$, so its off-path belief is unchanged. Hence, (ω, x) remains stable under ν'.

A.1.1 Proof of Lemma 2

Part (i). Suppose x is complete-information stable at ω and let $F = \{(\omega, x)\}$. For every firm j, set $\nu_j^0 = \delta_\omega$. For every deviation with a nonempty inference event, define

$$\nu_j^{(i,\hat{c})} = \begin{cases} \delta_\omega, & \text{if } \omega \in D_j(\omega, x; i, \hat{c}), \\ \text{Unif}(D_j(\omega, x; i, \hat{c})), & \text{otherwise.} \end{cases}$$

The resulting profile strongly believes $F_x = \{\omega\}$. Complete-information individual rationality implies individual rationality under ν. Every deviation accepted at ω is evaluated under δ_ω and therefore cannot strictly benefit the firm. Thus, (ω, x) is stable under ν, and F is self-stabilizing.

Conversely, suppose $F = \{(\omega, x)\}$ is self-stabilizing. Every supporting profile satisfies $\nu_j^0[\omega] = 1$ and $\nu_j^{(i,\hat{c})}[\omega] = 1$ whenever $\omega \in D_j(\omega, x; i, \hat{c})$. Stability under such a profile therefore implies individual rationality and no-blocking under complete information at ω.

Part (ii). Let F be the union of a nonempty collection of self-stabilizing sets. Fix $(\omega, x) \in F$ and choose a self-stabilizing set $\widehat{F}$ in the collection containing this outcome. By the monotonicity observation, a profile supporting (ω, x) in $\widehat{F}$ can be modified to strongly believe F_x while preserving stability. Since $(\omega, x) \in F$ was arbitrary, F is self-stabilizing.

A.1.2 Proof of Lemma 3

Fix $(\omega, x) \in F$, with $x = (\mu, \mathbf{c})$, and choose a conjecture profile $\widehat{\nu}$ supporting it in F. For each firm j, retain the on-path belief by setting $\nu_j^0 = \widehat{\nu}_j^0$. Consider any deviation $(i, j, \hat{c})$ with

a nonempty inference event. If $F_x \cap D_j(\omega, x; i, \hat{c}) = \varnothing$, retain $\widehat{\nu}_j^{(i,\hat{c})}$. Otherwise, choose

$$\omega' \in F_x \cap D_j(\omega, x; i, \hat{c})$$

and a conjecture profile supporting (ω', x) in F. Let $\nu_j^{(i,\hat{c})}$ be firm j's off-path belief for this deviation from that profile.

The selected belief is admissible at (ω, x) because the partitional structure implies

$$\Pi_j(\omega', x) = \Pi_j(\omega, x) \quad \text{and} \quad D_j(\omega', x; i, \hat{c}) = D_j(\omega, x; i, \hat{c}).$$

Strong belief at (ω', x) gives

$$\nu_j^{(i,\hat{c})}[F_x \cap D_j(\omega, x; i, \hat{c})] = 1. \tag{3}$$

Since worker i accepts the deviation at ω', stability of the selected profile also gives

$$\mathbb{E}_{\nu_j^{(i,\hat{c})}}[v_{ij}(\cdot, \hat{c})] \leq \mathbb{E}_{\nu_j^{(i,\hat{c})}}[v_{\mu(j)j}(\cdot, c_{\mu(j)j})]. \tag{4}$$

These beliefs can be selected independently across deviations and therefore define a conjecture ν_j.

The resulting profile strongly believes F_x and preserves individual rationality at (ω, x). Now take any $\omega'' \in F_x \cap \Pi_j(\omega, x)$. Every deviation accepted at ω'' has an inference event meeting F_x, so (4) rules out a strict gain to the firm. Hence,

$$F_x \cap \Pi_j(\omega, x) \subseteq NB_j(\omega, x; \nu_j).$$

In particular, no deviation blocks at ω, and (ω, x) is stable under ν.

Finally, the on-path belief assigns probability one to $F_x \cap \Pi_j(\omega, x)$ and therefore to $NB_j(\omega, x; \nu_j)$. Whenever an inference event meets F_x, its off-path belief also assigns probability one to this no-blocking set by (3). Thus, every ν_j is self-consistent relative to F_x.

A.1.3 Supporting a self-stabilizing set relative to a larger set

The next lemma allows us to apply the preceding result at each round of rationalizable α-stability.

Lemma 4. *Let $F \subseteq R \subseteq \mathcal{O}$, where F is self-stabilizing. For every $(\omega, x) \in F$, there exists a conjecture profile satisfying conditions (i)–(iii) of Definition 6, with R in place of R^{n-1}.*

Proof. Fix $(\omega, x) \in F$. By Lemma 3, choose a profile ν^F supporting (ω, x) in F such that

$$F_x \cap \Pi_j(\omega, x) \subseteq NB_j(\omega, x; \nu_j^F) \quad \text{for every } j \in J. \tag{5}$$

For each firm j, set $\nu_j^{(0)} = \nu_j^F$ and $R_j^0 = R_x \cap \Pi_j(\omega, x)$. At every step $\ell \geq 0$, keep the on-path belief unchanged and, for every deviation with a nonempty inference event, define

$$\nu_j^{(\ell+1),(i,\hat{c})} = \begin{cases} \text{Unif}(R_j^\ell \cap D_j(\omega, x; i, \hat{c})), & \text{if } F_x \cap D_j(\omega, x; i, \hat{c}) = \varnothing \\ & \text{and } R_j^\ell \cap D_j(\omega, x; i, \hat{c}) \neq \varnothing, \\ \nu_j^{(\ell),(i,\hat{c})}, & \text{otherwise.} \end{cases} \tag{6}$$

Let

$$R_j^{\ell+1} = R_x \cap NB_j(\omega, x; \nu_j^{(\ell+1)}).$$

Every update defines an admissible conjecture because each changed belief is supported on a nonempty subset of its inference event.

By construction, all off-path beliefs whose inference events meet F_x remain those of ν_j^F. Every deviation accepted at a state in $F_x \cap \Pi_j(\omega, x)$ has such an inference event. Therefore, (5) continues to hold after every update. Together with $F \subseteq R$, this implies

$$F_x \cap \Pi_j(\omega, x) \subseteq R_j^\ell \quad \text{for every } \ell \geq 0. \tag{7}$$

We next show that the sequence $(R_j^\ell)_{\ell \geq 0}$ is decreasing. The inclusion $R_j^1 \subseteq R_j^0$ is immediate. For $\ell \geq 1$, take any $\omega' \in R_x \cap \Pi_j(\omega, x)$ outside R_j^ℓ. Some deviation $(i, j, \hat{c})$ blocks (ω', x) under $\nu_j^{(\ell)}$. The same deviation blocks at every state in its inference event under that conjecture, so

$$R_j^\ell \cap D_j(\omega, x; i, \hat{c}) = \varnothing.$$

The update leaves the corresponding off-path belief unchanged. Thus, the deviation still blocks (ω', x) under $\nu_j^{(\ell+1)}$, and $\omega' \notin R_j^{\ell+1}$. Hence, $R_j^{\ell+1} \subseteq R_j^\ell$.

Since Ω is finite, choose $L_j \geq 1$ such that $R_j^{L_j} = R_j^{L_j - 1}$ and set $\nu_j = \nu_j^{(L_j)}$. We verify the required properties of $\nu = (\nu_j)_{j \in J}$.

Stability. Workers' individual rationality and firms' on-path beliefs are unchanged. Moreover, (7) and $\omega \in F_x$ imply $\omega \in NB_j(\omega, x; \nu_j)$ for every firm j. Thus, (ω, x) is stable under ν.

Strong belief. The on-path belief remains supported on $F_x \subseteq R_x$. At the first update, every off-path belief whose inference event meets R_x either remains supported on F_x or becomes uniform on $R_x \cap D_j(\omega, x; i, \hat{c})$. Hence, $\nu_j^{(1)}$ strongly believes R_x. Every subsequent update either preserves the preceding belief or replaces it with a distribution supported on $R_j^\ell \subseteq R_x$. Therefore, ν_j also strongly believes R_x.

Self-consistency. By (7), the unchanged on-path belief assigns probability one to $NB_j(\omega, x; \nu_j)$. Consider a deviation such that

$$R_x \cap NB_j(\omega, x; \nu_j) \cap D_j(\omega, x; i, \hat{c}) = R_j^{L_j} \cap D_j(\omega, x; i, \hat{c}) \neq \varnothing.$$

If its inference event meets F_x, the off-path belief remains supported on $F_x \cap \Pi_j(\omega, x) \subseteq R_j^{L_j}$. Otherwise, because $R_j^{L_j-1} = R_j^{L_j}$, the final update makes this belief uniform on $R_j^{L_j} \cap D_j(\omega, x; i, \hat{c})$. In either case,

$$\nu_j^{(i,\hat{c})}[NB_j(\omega, x; \nu_j)] \geq \nu_j^{(i,\hat{c})}[R_j^{L_j}] = 1.$$

Thus, every ν_j is self-consistent relative to R_x. □

A.1.4 Proof of Proposition 1

First step: The iteration without self-consistency. We first show that every self-stabilizing set F is contained in $\widehat{R}^\infty$. Since $\widehat{R}^0 = \mathcal{O}$, we have $F \subseteq \widehat{R}^0$. Suppose $F \subseteq \widehat{R}^{n-1}$. By the monotonicity observation, a profile supporting any $(\omega, x) \in F$ can be modified to strongly believe $\widehat{R}_x^{n-1}$ while preserving stability. Therefore, $(\omega, x) \in \widehat{R}^n$. Induction gives $F \subseteq \widehat{R}^\infty$.

Fix an allocation x. The sequence $(\widehat{R}_x^n)_{n\geq 0}$ is decreasing and starts from Ω. The test at (ω, x) depends on the preceding round only through $\widehat{R}_x^{n-1}$. Thus, if two consecutive sets are equal, all subsequent sets are equal. Otherwise, at least one state is removed. Since there can be at most $|\Omega|$ strict decreases,

$$\widehat{R}_x^{|\Omega|} = \widehat{R}_x^{|\Omega|+1} = \widehat{R}_x^\infty.$$

The bound is independent of x, so $\widehat{R}^{|\Omega|} = \widehat{R}^\infty$. The equality between consecutive rounds implies that every outcome in $\widehat{R}^\infty$ is stable under a profile that strongly believes $\widehat{R}_x^\infty$. Moreover, Lemmas 2(i) and 1, together with the preceding inclusion, imply that $\widehat{R}^\infty$ contains a complete-information stable outcome at every state. Hence, $\widehat{R}^\infty$ is the largest self-stabilizing set.

Second step: Comparison of the iterations. We show by induction that $R^n \subseteq \widehat{R}^n$ for every $n \geq 0$. The inclusion holds at $n = 0$. Suppose $R^{n-1} \subseteq \widehat{R}^{n-1}$ and take $(\omega, x) \in R^n$. Its supporting profile makes the outcome stable and strongly believes R_x^{n-1}. By the monotonicity observation, we can preserve stability while replacing this profile with one that strongly believes $\widehat{R}_x^{n-1}$. Since $(\omega, x) \in R^{n-1} \subseteq \widehat{R}^{n-1}$, it follows that $(\omega, x) \in \widehat{R}^n$.

For the reverse inclusion at the limit, let F be any self-stabilizing set. Since $R^0 = \mathcal{O}$, we have $F \subseteq R^0$. Suppose $F \subseteq R^{n-1}$. Applying Lemma 4 with R^{n-1} in place of R gives a conjecture profile supporting each outcome in F in round n. Thus, $F \subseteq R^n$ for every $n \geq 0$. Applying this conclusion to $F = \widehat{R}^\infty$ yields

$$\widehat{R}^\infty \subseteq R^\infty \subseteq \widehat{R}^\infty.$$

Therefore, the two limits coincide and equal the largest self-stabilizing set.

Finally,

$$R^\infty \subseteq R^{|\Omega|} \subseteq \widehat{R}^{|\Omega|} = \widehat{R}^\infty = R^\infty,$$

so both iterations terminate within $|\Omega|$ rounds. By Lemma 2(i), every complete-information stable outcome belongs to their common limit. Lemma 1 ensures that at least one such outcome exists at every state.

A.2 Proof of Proposition 2

(i) $\Rightarrow$ (ii). Fix $(\omega, x) \in R^\infty$, where $x = (\mu, \mathbf{c})$. We construct an uninformative expansion and an α-stable matching function that realizes (ω, x).

The expanded states and matching function. Let

$$\bar{\Omega} = (\Omega \times \{0\}) \cup (R_x^\infty \times \{1\}), \qquad \rho(\omega', s) = \omega',$$

and set $\bar{\tau} = \tau \circ \rho$. For every firm j, expanded state (ω', s), and allocation $y \in X$, define

$$\bar{\Pi}_j((\omega', s), y) = \rho^{-1}(\Pi_j(\omega', y)).$$

For each $\omega' \in \Omega$, choose a complete-information stable allocation $x^{\omega'}$, whose existence follows from Lemma 1. Define

$$M(\omega', 0) = x^{\omega'} \quad \text{for every } \omega' \in \Omega,$$
$$M(\omega', 1) = x \quad \text{for every } \omega' \in R_x^\infty.$$

If $x^{\omega'} = x$, Proposition 1(iv) implies $\omega' \in R_x^\infty$. Since x is assigned to every second copy, we obtain

$$\rho(M^{-1}(x)) = R_x^\infty. \tag{8}$$

In particular, $M(\omega, 1) = x$.

The expanded priors. For each $\omega' \in R_x^\infty$, choose a conjecture profile $\nu^{\omega'}$ supporting (ω', x) in R^∞ as in Lemma 3. Thus, for every firm j,

$$R_x^\infty \cap \Pi_j(\omega', x) \subseteq NB_j(\omega', x; \nu_j^{\omega'}).$$

Define

$$\pi_j = \frac{1}{|R_x^\infty|} \sum_{\omega' \in R_x^\infty} \nu_j^{\omega', 0}.$$

Each term is supported on $R_x^\infty \cap \Pi_j(\omega', x)$. Since the information sets form a partition, for every $\omega'' \in R_x^\infty$,

$$\pi_j[\Pi_j(\omega'', x)] = \frac{|R_x^\infty \cap \Pi_j(\omega'', x)|}{|R_x^\infty|} > 0. \tag{9}$$

Conditional on this information cell, π_j is the average of the supporting on-path beliefs at the states in $R_x^\infty \cap \Pi_j(\omega'', x)$. Each of these beliefs satisfies firm j's individual rationality, so

$$\mathbb{E}_{\pi_j}[v_{\mu(j)j}(\cdot, c_{\mu(j)j}) \,|\, \Pi_j(\omega'', x)] \geq 0. \tag{10}$$

Choose

$$0 < \varepsilon < \min_{j \in J, \omega' \in \Omega} \beta_j^0[\omega'],$$

which is possible because the original priors have full support. Define

$$\begin{aligned} \bar{\beta}_j^0[(\omega', 1)] &= \varepsilon \pi_j[\omega'] && \text{for } \omega' \in R_x^\infty, \\ \bar{\beta}_j^0[(\omega', 0)] &= \beta_j^0[\omega'] - \varepsilon \pi_j[\omega'] && \text{for } \omega' \in \Omega. \end{aligned} \tag{11}$$

Every first copy has strictly positive probability; second copies may have zero probability. Since π_j is supported on R_x^∞, for every $E \subseteq \Omega$,

$$\bar{\beta}_j^0[\rho^{-1}(E)] = \sum_{\omega' \in E} \left(\beta_j^0[\omega'] - \varepsilon \pi_j[\omega']\right) + \varepsilon \pi_j[E] = \beta_j^0[E].$$

Thus, $\bar{\beta}_j^0$ is a probability distribution preserving firm j's original prior. Together with the definitions of $\bar{\tau}$ and $\bar{\Pi}_j$, this establishes that $\bar{\mathcal{E}} = \langle \bar{\mathcal{I}}, \bar{\beta}^0 \rangle$ is an uninformative expansion of $\mathcal{E}$.

Stability at the target allocation. Take any $\bar{\omega} \in M^{-1}(x)$. By (8), $\rho(\bar{\omega}) \in R_x^\infty$, so workers' individual rationality follows from stability of the corresponding original outcome.

Fix a firm j. Her on-path conditioning event contains every second copy whose projection belongs to $\Pi_j(\rho(\bar{\omega}), x)$. Therefore, (9) implies

$$\bar{\beta}_j^0[M^{-1}(x) \cap \bar{\Pi}_j(\bar{\omega}, x)] \geq \varepsilon \pi_j[\Pi_j(\rho(\bar{\omega}), x)] > 0.$$

The contribution of these second copies to the numerator of her conditional expected payoff is

$$\varepsilon \sum_{\omega' \in \Pi_j(\rho(\bar{\omega}), x)} \pi_j[\omega'] v_{\mu(j)j}(\omega', c_{\mu(j)j}) \geq 0,$$

where the inequality follows from (10). Every first copy $(\omega', 0)$ in the conditioning event satisfies $x^{\omega'} = x$, so complete-information stability implies $v_{\mu(j)j}(\omega', c_{\mu(j)j}) \geq 0$. Adding these contributions and dividing by the positive probability of the conditioning event establishes firm j's individual rationality.

Now consider a deviation $(i, j, \hat{c})$ such that

$$M^{-1}(x) \cap \bar{D}_j(\bar{\omega}, x; i, \hat{c}) \neq \varnothing.$$

Payoff preservation and the definition of $\bar{\Pi}_j$ give

$$\bar{D}_j(\bar{\omega}, x; i, \hat{c}) = \rho^{-1}(D_j(\rho(\bar{\omega}), x; i, \hat{c})).$$

Together with (8), this implies

$$R_x^\infty \cap D_j(\rho(\bar{\omega}), x; i, \hat{c}) \neq \varnothing.$$

By the choice of $\nu_j^{\rho(\bar{\omega})}$, its off-path belief for this deviation is supported on this intersection and rules out blocking at every state in it. Since the worker accepts throughout the intersection,

$$\mathbb{E}_{\nu_j^{\rho(\bar{\omega}),(i,\hat{c})}}[v_{ij}(\cdot, \hat{c})] \leq \mathbb{E}_{\nu_j^{\rho(\bar{\omega}),(i,\hat{c})}}[v_{\mu(j)j}(\cdot, c_{\mu(j)j})].$$

Lift this belief to the second copies by assigning probability $\nu_j^{\rho(\bar{\omega}),(i,\hat{c})}[\omega']$ to each $(\omega', 1)$ and zero to every first copy. The lifted belief is supported on $M^{-1}(x) \cap \bar{D}_j(\bar{\omega}, x; i, \hat{c})$ and preserves both expected payoffs. It therefore satisfies condition (iii) of Definition 8. These choices can be made separately for each information cell and deviation.

Stability at the other allocations. Take any $y \neq x$ in the range of M. Every state inducing y is a first copy $(\omega', 0)$ with $x^{\omega'} = y$. Each such copy has strictly positive probability under every firm's prior, so every on-path conditioning event has positive probability. Complete-information stability at these states gives workers' individual rationality and nonnegative firm payoffs throughout the support of each on-path posterior.

For any deviation whose inference event intersects $M^{-1}(y)$, the worker accepts at every state in the intersection. Complete-information stability implies that the firm does not strictly gain at any of those states. Hence, any belief supported on the intersection satisfies condition (iii) of Definition 8. Thus, M is α-stable on $\bar{\mathcal{E}}$ and realizes (ω, x) at $(\omega, 1)$.

(ii) $\Rightarrow$ (iii). Every uninformative expansion is an expansion.

(iii) $\Rightarrow$ (i). Let $\bar{\mathcal{E}}$ be an expansion of $\mathcal{E}$ with projection ρ, and let M be an α-stable matching function on $\bar{\mathcal{E}}$. Define

$$F = \{(\rho(\bar{\omega}), M(\bar{\omega})) : \bar{\omega} \in \bar{\Omega}\}.$$

In particular, $F_x = \rho(M^{-1}(x))$. We show that F is self-stabilizing.

Fix $(\omega, x) \in F$, with $x = (\mu, \mathbf{c})$, and choose $\bar{\omega} \in \rho^{-1}(\omega)$ such that $M(\bar{\omega}) = x$. For every firm j and $E \subseteq \Omega$, define

$$\nu_j^0[E] = \bar{\beta}_j^0[\rho^{-1}(E) \mid M^{-1}(x) \cap \bar{\Pi}_j(\bar{\omega}, x)].$$

Condition (ii) of Definition 8 ensures that the conditioning event has positive probability. Information refinement implies that ν_j^0 is supported on $F_x \cap \Pi_j(\omega, x)$. Moreover, payoff preservation gives

$$\begin{aligned}\mathbb{E}_{\nu_j^0}[v_{\mu(j)j}(\cdot, c_{\mu(j)j})] &= \mathbb{E}_{\bar{\beta}_j^0}[v_{\mu(j)j}(\rho(\cdot), c_{\mu(j)j}) \mid M^{-1}(x) \cap \bar{\Pi}_j(\bar{\omega}, x)] \\ &\geq 0.\end{aligned}$$

Workers' individual rationality also carries over from $(\bar{\omega}, x)$.

For each deviation $(i, j, \hat{c})$ such that $F_x \cap D_j(\omega, x; i, \hat{c}) \neq \varnothing$, choose ω' in this intersection and $\bar{\omega}' \in \rho^{-1}(\omega')$ with $M(\bar{\omega}') = x$. The worker accepts the deviation at $\bar{\omega}'$. Hence, condition (iii) of Definition 8 provides a belief $\bar{\beta}_j^{(i,\hat{c})}$ supported on $M^{-1}(x) \cap \bar{D}_j(\bar{\omega}', x; i, \hat{c})$ under which the firm does not gain. Since $\omega' \in \Pi_j(\omega, x)$, information refinement and the partitional structure imply

$$\rho\Big(M^{-1}(x) \cap \bar{D}_j(\bar{\omega}', x; i, \hat{c})\Big) \subseteq F_x \cap D_j(\omega, x; i, \hat{c}).$$

For every $E \subseteq \Omega$, define

$$\nu_j^{(i,\hat{c})}[E] = \bar{\beta}_j^{(i,\hat{c})}[\rho^{-1}(E)].$$

This belief is supported on $F_x \cap D_j(\omega, x; i, \hat{c})$ and satisfies

$$\mathbb{E}_{\nu_j^{(i,\hat{c})}}[v_{ij}(\cdot, \hat{c})] \leq \mathbb{E}_{\nu_j^{(i,\hat{c})}}[v_{\mu(j)j}(\cdot, c_{\mu(j)j})].$$

For every remaining deviation with a nonempty inference event, choose any belief supported on that event. The selections can be made independently across deviations.

The resulting profile strongly believes F_x. Every deviation accepted at ω has an inference event meeting F_x, because $\omega \in F_x$, and is deterred by the corresponding off-path belief. Together with individual rationality, this establishes stability of (ω, x) under ν. Since $(\omega, x) \in F$ was arbitrary, F is self-stabilizing. Proposition 1(ii) then gives $F \subseteq R^\infty$, completing the proof.

References

Battigalli, Pierpaolo, and Marciano Siniscalchi (2002): "Strong Belief and Forward Induction Reasoning," *Journal of Economic Theory*, 106(2), 356–391.

Battigalli, Pierpaolo, and Marciano Siniscalchi (2003): "Rationalization and Incomplete Information," *Advances in Theoretical Economics*, 3(1), Article 3.

Bergemann, Dirk, and Stephen Morris (2017): "Belief-free Rationalizability and Informational Robustness," *Games and Economic Behavior*, 104, 744–759.

Bernheim, B. Douglas (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52(4), 1007–1028.

Brandenburger, Adam, and Eddie Dekel (1987): "Rationalizability and Correlated Equilibria," *Econometrica*, 55(6), 1391–1402.

Fisher, James C. D. (2020): "Existence of Stable Allocations in Matching Markets with Infinite Contracts: A Topological Approach," *Journal of Mathematical Economics*, 91, 136–140.

Gale, David, and Lloyd S. Shapley (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69(1), 9–15.

Hatfield, John William, and Paul R. Milgrom (2005): "Matching with Contracts," *American Economic Review*, 95(4), 913–935.

Kotowski, Maciej H. (2024): "A Perfectly Robust Approach to Multiperiod Matching Problems," *Journal of Economic Theory*, 222, 105919.

Liu, Qingmin (2020): "Stability and Bayesian Consistency in Two-Sided Markets," *American Economic Review*, 110(8), 2625–2666.

Liu, Qingmin, George J. Mailath, Andrew Postlewaite, and Larry Samuelson (2014): "Stable Matching with Incomplete Information," *Econometrica*, 82(2), 541–587.

Pearce, David G. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52(4), 1029–1050.

Pomatto, Luciano (2022): "Stable Matching under Forward-Induction Reasoning," *Theoretical Economics*, 17(4), 1619–1649.

Shapley, Lloyd S., and Martin Shubik (1971): "The Assignment Game I: The Core," *International Journal of Game Theory*, 1, 111–130.

Stalnaker, Robert (1998): "Belief Revision in Games: Forward and Backward Induction," *Mathematical Social Sciences*, 36(1), 31–56.

Wang, Ziwei (2026): "Rationalizable Stability in Matching with Incomplete Information," Working paper, Peking University.

Nicolò, Antonio, Riccardo Saulle, and Pietro Salmaso (2026): "Rationalizable Conjectures in Dynamic Matching," Working paper.

Nicolò, Antonio, Pietro Salmaso, and Riccardo Saulle (2026): "Rationalizable Behavior in Matching with Externalities," Working paper.

Penta, Antonio (2015): "Robust Dynamic Implementation," *Journal of Economic Theory*, 160, 280–316.

Catonini, Emiliano and Antonio Penta (2026): "Backward Induction Reasoning beyond Backward Induction," *American Economic Journal: Microeconomics*, 18(1), 30–58.

Fudenberg, Drew and Jean Tirole (1991): *Game Theory*. Cambridge, MA: MIT Press.