False Hope and the Paradox of Plenty: Overconfidence Breaks Matching When Search is Cheap

Justin Tumlinson* Felix Várdy†

<PRELIMINARY AND INCOMPLETE>

Abstract

We show that even small belief distortions can destroy matching when search is cheap. In a two-sided market with sequential search, we introduce a small wedge between self-worth and market value. When agents are unaware of this wedge, cheap search generates a paradox of plenty: away from the terminal layer, mutual acceptance collapses on any compact bounded away from zero, yet agents continue to search because they perceive nonempty opportunities ("false hope"). Under partial awareness—agents recognize others' delusions but not their own—some actual overlap survives on certain compacts, but cheap search still creates substantial regions of false hope: as the number of draws grows, the mass of agents who perceive mutually acceptable partners while in fact none exist becomes dominant. Under full awareness, where wedges are common knowledge, the geometry of acceptance is sharply determined by primitives.

1 Introduction

Search is often treated as an unambiguous remedy for mismatch: more opportunities, lower frictions, better outcomes. This paper shows that the logic can reverse when agents are even slightly wrong about their market value or attractiveness. When search becomes arbitrarily cheap, small belief distortions can undermine—indeed, in some environments destroy—mutual acceptance away from the end of the horizon.

*University of Exeter. Email: J.Tumlinson@exeter.ac.uk.

†International Monetary Fund. Email: fvardy@gmail.com.

We study a two-sided matching market with sequential search over a long horizon. Agents observe partners' market values, but evaluate relationships through a wedge between market value and perceived self-worth. The wedge is intentionally modest; its force comes from compounding through best responses over many search opportunities. Our focus is not on exotic preferences or extreme frictions, but on how a thin layer of misperception interacts with near-frictionless search.

The key distinction is informational. We compare three environments that differ only in what agents understand about the wedge. Under unawareness, agents treat perceived self-worth as market value. In this case, arbitrarily cheap search produces a Paradox of Plenty: away from the terminal layer, mutual acceptance collapses on any compact bounded away from zero. Yet search does not stop. Agents continue to search because they perceive non-empty opportunities—what we call false hope—even though mutual acceptance is essentially absent.

Under partial awareness, agents recognize others' delusions but not their own. Here the stark collapse is generically avoided: some actual overlap survives on compacts. But cheap search still produces large regions of false hope. As the number of draws grows, the mass of agents who believe they have mutually acceptable partners while in fact none exist becomes dominant on interior layers. In other words, the paradox shifts form: matching does not vanish uniformly, but futile search becomes the typical experience.

Under full awareness, wedges are common knowledge. In that environment, acceptance geometry is sharply pinned down by primitives. Cheap search does not generate the same amplification of mistaken beliefs, and the model delivers crisp comparative statics that contrast with the unawareness and partial-awareness regimes.

Our results speak to a broader tension in modern matching environments—dating markets, hiring platforms, and other settings where search costs are low and feedback is noisy. Cheap search expands opportunity, but it also magnifies the consequences of small, systematic errors in self-assessment and inference. The paper formalizes this mechanism, characterizes when it leads to breakdown versus segmentation, and quantifies how quickly false hope absorbs the population as search becomes cheap.

2 Model

2.1 Setup

Consider a two-sided matching market that unfolds over $n \in \mathbb{N}$ macro periods. Each macro period is subdivided into $m \in \mathbb{N}$ sequential sub-periods ("draws"), so the effective horizon is $N \equiv nm$. Let $H_i \equiv 1 + N - i$ denote the number of sub-periods remaining, including sub-period i.

Interpret each macro period as one unit of time and each draw as having length $1/m$. Flow payoffs (single or matched) are naturally measured per macro period. We therefore normalize payoffs and value functions by draw length (i.e., multiply by m), so that per-draw payoffs coincide with the corresponding flow payoffs. This normalization does not affect acceptance decisions or equilibrium cutoffs; to recover unnormalized utilities, divide value functions by m.

In each sub-period $i \in \{1, \ldots, N\}$, an agent on side $k \in \{A, B\}$ with market value $t_k \geq 0$ meets a potential partner on side $l \in \{A, B\}$, $l \neq k$, with observable market value $t_l \geq 0$ drawn i.i.d. from a common distribution F with density f that is continuous on $(0, \infty)$, and finite expectation. Furthermore, $f[t] > 0$ for all $t > 0$. An agent with market value t_k has self-worth $\lambda_k t_k$, where $\lambda_k > 0$ parametrizes the wedge between market value and self-worth for side $k \in \{A, B\}$. A value $\lambda_k > 1$ means that self-worth exceeds market value, while $\lambda_k < 1$ means that self-worth is less than market value. Let s_k denote an agent's self-worth (so $s_k = \lambda_k t_k$).

We consider three levels of awareness:

1. **Unawareness (UA):** Agents are unaware of the wedges between market value and self-worth. They behave as if $\lambda_A = \lambda_B = 1$: they interpret their own self-worth as their market value, and they interpret their potential partner's observed market value as the partner's self-worth.

2. **Full Awareness (FA):** Agents are fully aware of the wedges between market value and self-worth. They understand that their market value is t_k while their self-worth is $\lambda_k t_k$, and this is common knowledge.

3. **Partial Awareness (PA):** Agents recognize that their potential partners treat self-worth as market value, but they do not realize that they themselves are making the

same mistake. That is, they wrongly believe their market value to be s_k.

After observing their partner's market value, the pair decide whether to form a couple. They exit the market if and only if both consent. In that case, agent A with market value t_A receives a flow payoff of $\kappa(s_A + t_B)/2$, where $1 < \kappa < 2$, while agent B with market value t_B receives $\kappa(t_A + s_B)/2$. If either agent refuses, both remain single in this draw, collect a normalized stage payoff equal to their self-worth (s_A and s_B, respectively), and enter the next draw, where the matching process repeats. Future payoffs are discounted by $\delta \in (0, 1)$ per macro period. Thus, the per-draw discount factor is $\hat{\delta} \equiv \delta^{1/m}$.

3 Equilibrium

In FA, the model is common knowledge, and the appropriate equilibrium concept is subgame perfect equilibrium (SPE). We ignore SPEs in weakly dominated strategies, in which agents reject solely because they expect their potential partner to reject. In UA and PA, the model is *not* common knowledge, with agents misperceiving reality. In those circumstances, we solve for *subjective* subgame perfect equilibrium (SSPE), in which agents best-respond to their (incorrect) subjective beliefs, again ignoring equilibria in weakly dominated strategies. This equilibrium is subjective because the dynamic program is solved under these mistaken beliefs. Once the subjective cutoff functions are pinned down, we can translate them into the actual mutual-acceptance set by substituting each side's self-perception $s_k = \lambda_k t_k$ into the cutoff rules.

3.1 UA

First, suppose agents are wholly unaware of their own self-delusion and of their partner's self-delusion. An agent with true type t_k self-perceives as $s_k = \lambda_k t_k$, but—being unaware of this wedge—treats s_k as his "true" type when forming beliefs and choosing acceptance. Likewise, he treats the partner's type t_l as the partner's true type. As a consequence, agents (incorrectly) view the environment as symmetric, even when $\lambda_A \neq \lambda_B$.

Let $\rho_{k,i}[s_k]$ denote the cutoff rule that an agent of type t_k who self-perceives as s_k applies in sub-period i : after observing the partner's (true) type t_l, he accepts iff

$$t_l \geq \rho_{k,i}\left[s_k\right].$$

Not accounting for the partner's self-delusion, he believes the partner accepts him whenever $s_k \geq \rho_{l,i}[t_l]$. Because $\rho_{l,i}[\cdot]$ is strictly increasing (Lemma 11), this is equivalent to $t_l \leq \rho_{l,i}^{-1}[s_k]$. Here, $\rho_{k,i}^{-1}$ denotes the generalized inverse of $\rho_{k,i}$, defined for $t \in \mathbb{R}_+$ by

$$\rho_{k,i}^{-1}[t] \equiv \inf\{s \geq 0 : \rho_{k,i}[s] \geq t\}.$$

This generalized inverse coincides with the usual functional inverse on Range $\left[\rho_{k,i}\right]$ and satisfies $\rho_{k,i}^{-1}[\rho_{k,i}[s]] = s$ for all $s \geq 0$ (and $\rho_{k,i}[\rho_{k,i}^{-1}[t]] = t$ for all $t \in$ Range $\left[\rho_{k,i}\right]$).

Agent k (falsely) believes that a match forms in sub-period i iff

$$t_l \in \left[\rho_{k,i}\left[s_k\right], \rho_{l,i}^{-1}\left[s_k\right]\right]. \tag{1}$$

The lower endpoint of this interval is simply k's own acceptance constraint, while the upper endpoint reflects the highest partner type that k expects would accept him given his perceived type s_k. Thus, under UA, agent k treats the mutual-acceptance requirement as a "band" around s_k: partners that are too low are rejected by k, while partners, l, that are too high reject k. The perceived upper bound, $\rho_{l,i}^{-1}\left[s_k\right]$, will play a role below when we select a canonical representative among payoff-equivalent cutoffs $\rho_{k,i}\left[s_k\right]$.

3.2 FA

Next, suppose agents are fully aware of the wedges between market value and self-worth, and this is common knowledge. Thus, a side-k agent with market value t_k correctly understands that his self-worth is $s_k = \lambda_k t_k$, and that a side-l opponent with market value t_l has self-worth $s_l = \lambda_l t_l$. Relative to UA, the primitive payoffs are unchanged; what changes is purely informational. In particular, there is no longer a discrepancy between the perceived and actual mutual acceptance regions.

Two features are worth keeping in mind. First, even under FA, self-worth still enters payoffs: when single, the agent receives s_k per draw, while a match yields the flow payoff $\kappa(s_k + t_l)/2$ for side k (and $\kappa(t_k + s_l)/2$ for side l). Hence, while a larger wedge λ_k raises both payoffs, the payoff from being single rises faster than the payoff from matching, making matching harder to justify. Second, when $\lambda_A \neq \lambda_B$, the environment is genuinely asymmetric and known to be so. Unlike in the UA case, the two sides do not generally share the same continuation values or acceptance thresholds.

Fix a subperiod i and consider a side-k agent with type t_k and self-worth $s_k = \lambda_k t_k$. As before, the optimal acceptance rule is of cutoff form: there exists a threshold function $\rho_{k,i}[\cdot]$ such that the agent accepts a partner of market value t_l iff

$$t_l \geq \rho_{k,i}[s_k].$$

The interpretation is the same as under UA, except that s_k is now correctly understood as the agent's self-worth, not as a misperceived market value. Under FA, the agent also correctly accounts for the opponent's acceptance decision. The side-l partner, with self-worth $s_l = \lambda_l t_l$, accepts the side-k agent iff

$$t_k \geq \rho_{l,i}[s_l] = \rho_{l,i}[\lambda_l t_l].$$

Since $\rho_{l,i}[\cdot]$ is strictly increasing (Lemma 11), this can equivalently be written as an upper bound on t_l,

$$t_l \leq \frac{1}{\lambda_l}\, \rho_{l,i}^{-1}[t_k].$$

Hence, conditional on t_k, the set of partner types that lead to mutual acceptance in subperiod i is the interval

$$t_l \in \left[\rho_{k,i}[\lambda_k t_k],\ \frac{1}{\lambda_l}\rho_{l,i}^{-1}[t_k]\right]. \tag{2}$$

This is the FA analogue of (1) under UA, with the key difference that here these bounds are not merely subjectively perceived; they are the actual mutual-acceptance constraints.

3.3 Partial Awareness

Finally, suppose that agents correctly understand that others misperceive their own types, but as in UA they themselves treat their self-worth $s_k = \lambda_k t_k$ as their market value. In addition, each agent believes that others *share* his incorrect self-belief about him. Concretely: (i) an agent with market value t_k thinks his value is $s_k = \lambda_k t_k$ and believes others view him as s_k; (ii) he is aware that a partner of market value t_l perceives herself as $s_l = \lambda_l t_l$.

An agent with self-perceived type $s_k = \lambda_k t_k$ is willing to partner with t_l in subperiod i iff $t_l \geq \rho_{k,i}[s_k]$. At the same time, he predicts that he is acceptable to a partner t_l whenever his (self-perceived) type clears the partner's cutoff evaluated at the partner's self-perceived type, s_l, i.e., iff

$$s_k \geq \rho_{l,i}[s_l] = \rho_{l,i}[\lambda_l t_l].$$

Equivalently,

$$\lambda_k t_k \geq \rho_{l,i}[\lambda_l t_l] \quad \Longleftrightarrow \quad t_l \leq \frac{1}{\lambda_l}\rho_{l,i}^{-1}[\lambda_k t_k].$$

Hence, under PA, type t_k *believes* that a match forms in subperiod i with partner type t_l iff

$$t_l \in \left[\rho_{k,i}[s_k],\ \frac{1}{\lambda_l}\rho_{l,i}^{-1}[s_k]\right], \tag{3}$$

which differs from UA only through the factor $1/\lambda_l$ in front of the perceived upper bound.

3.4 (Subjective) SPE

Using dynamic programming, we now show that, in all three cases, there exists an essentially unique (S)SPE.

Proposition 1 *Equilibrium exists in all three cases, UA, FA, PA. The associated (subjective) value functions are unique, and the cutoff rules are unique up to payoff-irrelevant boundary modifications. The actual mutual acceptance region is*

$$M_i = \left\{(t_A, t_B) \in \mathbb{R}_+^2 : \rho_{A,i}\left[\lambda_A t_A\right] \leq t_B \leq \frac{1}{\lambda_B}\rho_{B,i}^{-1}\left[t_A\right]\right\}\ .$$

1. ***UA:*** *In all periods, cutoff and value functions are the same for A and B. Type t_k with self-perception $s_k = \lambda_k t_k$ (falsely) believes that a match forms in i with t_l iff $t_l \in \left[\rho_{k,i}\left[s_k\right], \rho_{l,i}^{-1}\left[s_k\right]\right]$.*

2. ***FA:*** *Cutoff and value functions are generically side-specific. Perceived and actual mutual-acceptance regions coincide.*

3. ***PA:*** *Cutoff and value functions are generically side-specific. Type t_k with self-perception $s_k = \lambda_k t_k$ (falsely)believes that a match forms in i with t_l iff $t_l \in \left[\rho_{k,i}\left[s_k\right], \frac{1}{\lambda_l}\rho_{l,i}^{-1}\left[s_k\right]\right]$.*

Proof. We proceed by backward induction on subperiods. Fix a side $k \in \{A, B\}$ and suppress payoff-irrelevant tie-breaking on sets of partner types that occur with probability zero.

Case 1: UA

Under UA, both sides face the same subjective environment; hence it suffices to characterize the equilibrium objects for a generic side k, and symmetry then implies $\rho_{A,i}[s] = \rho_{B,i}[s]$ and $V_{A,i}[s] = V_{B,i}[s]$.

Fix a subperiod i and perceived type s_k. If a match forms in subperiod i with partner type t_l, the agent expects to receive the flow payoff $\frac{\kappa}{2}(s_k + t_l)$ in every remaining subperiod. The present value of matching from i is therefore

$$\alpha_i \frac{\kappa}{2}(s_k + t_l), \qquad \alpha_i \equiv \sum_{h=0}^{H_i - 1} \hat{\delta}^h = \frac{1 - \hat{\delta}^{H_i}}{1 - \hat{\delta}}.$$

If he rejects or is rejected, his continuation payoff is $s_k + \hat{\delta} V_{k,i+1}[s_k]$, with terminal condition $V_{k,N+1}[s] = 0$. Hence the cutoff $\rho_{k,i}[s_k]$ is characterized by indifference at the boundary partner type:

$$\alpha_i \frac{\kappa}{2} \left(s_k + \rho_{k,i}[s_k] \right) = s_k + \hat{\delta} V_{k,i+1}[s_k],$$

which yields the explicit threshold formula

$$\rho_{k,i}[s_k] = \frac{1 - \hat{\delta}}{1 - \hat{\delta}^{H_i}} \frac{2}{\kappa} \left(s_k + \hat{\delta} V_{k,i+1}[s_k] \right) - s_k. \tag{4}$$

In particular, in the final subperiod $i = N$ (where $\alpha_N = 1$),

$$\rho_{k,N}[s] = \left(\frac{2}{\kappa} - 1 \right) s,$$

so $\rho_{A,N}[\cdot] = \rho_{B,N}[\cdot] \equiv \rho_N[\cdot]$.

Given $\rho_{k,i}[\cdot]$, agent k (falsely) believes that a match forms in subperiod i precisely when

$$t_l \in \left[\rho_{k,i}[s_k], \rho_{l,i}^{-1}[s_k] \right].$$

Let

$$\pi_i[s] \equiv \max\{0,\ F(\rho_i^{-1}[s]) - F(\rho_i[s])\}$$

denote the perceived match probability, and interpret $\int_a^b t f(t)\, dt$ as 0 whenever $b \leq a$. Then the perceived value function satisfies

$$V_{k,i}[s_k] = (1 - \pi_i[s_k]) \left(s_k + \hat{\delta} V_{k,i+1}[s_k] \right) + \alpha_i \frac{\kappa}{2} \left(\pi_i[s_k]\, s_k + \int_{\rho_i[s_k]}^{\rho_i^{-1}[s_k]} t f(t)\, dt \right). \tag{5}$$

Starting from $V_{k,N+1} = 0$, equations (4) and (5) define $(\rho_{k,i}, V_{k,i})$ by backward recursion. The value functions are unique, and the cutoff rules are unique up to payoff-irrelevant modifications on probability-zero boundaries.

Finally, because under UA both sides solve the same recursion with the same terminal condition, it follows by backward induction that $\rho_{A,i}[s] = \rho_{B,i}[s] \equiv \rho_i[s]$ and $V_{A,i}[s] = V_{B,i}[s] \equiv V_i[s]$ for all i and s.

Matching region. A type-t_k agent self-perceives as $s_k = \lambda_k t_k$ and accepts t_l iff $t_l \geq \rho_i[\lambda_k t_k]$. He is accepted by t_l iff $t_k \geq \rho_i[\lambda_l t_l]$, equivalently $t_l \leq \frac{1}{\lambda_l}\rho_i^{-1}[t_k]$. Hence the *actual* mutual acceptance region in subperiod i is

$$M_i = \left\{(t_A, t_B) \in \mathbb{R}_+^2 : \ \rho_i[\lambda_A t_A] \leq t_B \leq \tfrac{1}{\lambda_B}\rho_i^{-1}[t_A]\right\}.$$

Equivalently,

$$M_i = \left\{(t_A, t_B) \in \mathbb{R}_+^2 : \ \rho_i[\lambda_B t_B] \leq t_A \leq \tfrac{1}{\lambda_A}\rho_i^{-1}[t_B]\right\},$$

since $\rho_i[\lambda_A t_A] \leq t_B \iff t_A \leq \frac{1}{\lambda_A}\rho_i^{-1}[t_B]$ and $t_B \leq \frac{1}{\lambda_B}\rho_i^{-1}[t_A] \iff \rho_i[\lambda_B t_B] \leq t_A$.

Case 2: FA

Step 1 (terminal subperiod). In subperiod $i = N$, there is no continuation value, since $V_{k,N+1} = 0$. Conditional on meeting a partner of market value t_l, accepting yields the match value $\frac{\kappa}{2}(s_k + t_l)$, while rejecting yields the single payoff s_k. Thus the optimal rule is of cutoff form and is uniquely pinned down by indifference at the boundary:

$$\frac{\kappa}{2}\left(s_k + \rho_{k,N}[s_k]\right) = s_k \qquad \Longrightarrow \qquad \rho_{k,N}[s] = \left(\frac{2}{\kappa} - 1\right)s.$$

Given these terminal cutoffs, $V_{k,N}[\cdot]$ is uniquely determined by the law of motion of the matching process and the (true) mutual-acceptance event.

Step 2 (inductive step). Fix $i \leq N-1$ and assume $V_{A,i+1}[\cdot]$ and $V_{B,i+1}[\cdot]$ are uniquely pinned down. Consider a side-k agent with self-worth s_k who meets a partner of type t_l in subperiod i. If both accept, the agent exits and obtains the present value $\alpha_i \frac{\kappa}{2}(s_k + t_l)$. If either party rejects, no match forms and the agent receives $s_k + \hat{\delta} V_{k,i+1}[s_k]$. Because the match value is strictly increasing in t_l, the set of partner types for which accepting is (weakly) optimal is an upper interval. The cutoff is therefore characterized by indifference at the boundary partner type,

$$\alpha_i \frac{\kappa}{2}\left(s_k + \rho_{k,i}[s_k]\right) = s_k + \hat{\delta} V_{k,i+1}[s_k],$$

which has a unique solution for $\rho_{k,i}[s_k]$. Hence the optimal strategy in subperiod i is a (side-specific) cutoff rule $t_l \geq \rho_{k,i}[s_k]$.

Given the pair of cutoff rules $(\rho_{A,i}, \rho_{B,i})$ and the type distribution, the value functions $V_{A,i}[\cdot]$ and $V_{B,i}[\cdot]$ are then uniquely determined by taking expectations over the mutually acceptable pairs in subperiod i, with continuation value $s_k + \hat{\delta} V_{k,i+1}[s_k]$ whenever no match

forms. This completes the backward recursion and yields a pure-strategy subgame-perfect equilibrium.

Step 3 (essential uniqueness and the matching region). Uniqueness of the value functions follows because the backward recursion pins down $V_{k,i}$ uniquely at each step given the already fixed future objects. The cutoff rules are unique up to payoff-irrelevant modifications on the boundary where the agent is indifferent. Finally, in equilibrium a match forms in subperiod i for (t_A, t_B) iff each side's cutoff constraint holds, i.e. iff $t_B \geq \rho_{A,i}[\lambda_A t_A]$ and $t_A \geq \rho_{B,i}[\lambda_B t_B]$, equivalently

$$t_B \in \left[\rho_{A,i}[\lambda_A t_A],\ \frac{1}{\lambda_B}\rho_{B,i}^{-1}[t_A]\right],$$

which yields the stated expression for M_i.

Case 3 (Partial Awareness):

Step 1 (terminal subperiod). Since $V_{k,N+1}[s] = 0$, rejecting in subperiod $i = N$ yields the single payoff s while accepting a partner of type t_l yields the match value $\frac{\kappa}{2}(s+t_l)$. Hence the best response is a cutoff rule, and the terminal cutoff is uniquely pinned down by indifference:

$$\frac{\kappa}{2}\left(s + \rho_{k,N}[s]\right) = s \qquad \Longrightarrow \qquad \rho_{k,N}[s] = \left(\frac{2}{\kappa} - 1\right) s.$$

This pins down $\rho_{k,N}$ uniquely (up to boundary modifications), and given $(\rho_{A,N}, \rho_{B,N})$ the subjective value functions $V_{A,N}[\cdot]$ and $V_{B,N}[\cdot]$ are uniquely determined from the subjective law of motion: a match is believed to form at (N, s_k) precisely when $t_l \in [\rho_{k,N}[s_k],\ \frac{1}{\lambda_l}\rho_{l,N}^{-1}[s_k]]$.

Step 2 (inductive step). Fix $i \leq N-1$ and suppose $V_{A,i+1}[\cdot]$ and $V_{B,i+1}[\cdot]$ have been uniquely pinned down. Consider a side-k agent with perceived type s who meets a partner of market value t_l in subperiod i. If both accepts, he exits and obtains $\alpha_i \frac{\kappa}{2}(s+t_l)$. If no match forms, he receives the continuation payoff $s + \hat{\delta} V_{k,i+1}[s]$. Because the match value is strictly increasing in t_l, the set of partner types for which accepting is (weakly) optimal is an upper interval. The cutoff therefore solves the indifference condition at the boundary partner type,

$$\alpha_i \frac{\kappa}{2}\left(s + \rho_{k,i}[s]\right) = s + \hat{\delta} V_{k,i+1}[s],$$

which has a unique solution for $\rho_{k,i}[s]$. Given the pair of cutoff rules $(\rho_{A,i}, \rho_{B,i})$, the subjective value functions $V_{A,i}[\cdot]$ and $V_{B,i}[\cdot]$ are uniquely determined by taking expectations over the *subjectively* mutually acceptable pairs in subperiod i, with continuation value $s + \hat{\delta} V_{k,i+1}[s]$ whenever no match is believed to form. A match is believed to form at (i, s_k) when $t_l \in$

$[\rho_{k,i}[s_k],\ \frac{1}{\lambda_l}\rho_{l,i}^{-1}[s_k]]$. This completes the backward recursion and establishes existence of a pure-strategy subjective SPE and essential uniqueness of $(V_{k,i}, \rho_{k,i})$ at each subperiod.

Step 3 (generic asymmetry). Unless $\lambda_A = \lambda_B$, the two subjective recursions do not coincide: each side's perceived upper-acceptance event involves the factor $1/\lambda_l$, so the subjective match probabilities and expected partner values generally differ across sides. Consequently, except on knife-edge primitives where the relevant continuation values coincide, we have $\rho_{A,i}[s] \neq \rho_{B,i}[s]$ and $V_{A,i}[s] \neq V_{B,i}[s]$ for $i \leq N-1$.

Step 4 (actual acceptance region). In reality, a side-A agent of market value t_A (with self-worth $s_A = \lambda_A t_A$) accepts in subperiod i iff $t_B \geq \rho_{A,i}[s_A] = \rho_{A,i}[\lambda_A t_A]$, and a side-$B$ agent of market value t_B accepts iff $t_A \geq \rho_{B,i}[\lambda_B t_B]$. Using strict monotonicity of $\rho_{B,i}[\cdot]$ (Lemma 11) and inverting the latter inequality yields $t_B \leq \frac{1}{\lambda_B}\rho_{B,i}^{-1}[t_A]$. Thus a match forms at (t_A, t_B) in subperiod i iff

$$\rho_{A,i}[\lambda_A t_A] \leq t_B \leq \frac{1}{\lambda_B}\rho_{B,i}^{-1}[t_A],$$

which is exactly the stated expression for M_i. ■

When agents are unaware of the wedges between market value and self-perception, both sides face the same perceived problem. Hence, the equilibrium cutoff and value functions coincide across sides. These subjective objects are then mapped into the actual matching region M_i in the space of true types (t_A, t_B), using the fact that each side applies the common cutoff to its self-perceived type $\lambda_k t_k$. In words, the set M_i is obtained by intersecting two one-sided constraints. The lower bound $t_B \geq \rho_i[\lambda_A t_A]$ says that agent A (who self-perceives as $\lambda_A t_A$) is willing to accept B. The upper bound $t_B \leq \frac{1}{\lambda_B}\rho_i^{-1}[t_A]$ is equivalent to $t_A \geq \rho_i[\lambda_B t_B]$ and therefore ensures that B is willing to accept A. This asymmetry between the two bounds is the key wedge created by self-delusion.

FA replaces the single-recursion structure of UA with a coupled system: each side's continuation values depend on the opponent's cutoffs via the expected quality of feasible matches. Nevertheless, the same backward-induction logic applies. In the terminal subperiod $i = N$, there is no continuation value, and each side's threshold reduces to the static one,

$$\rho_{k,N}[s] = \left(\frac{2}{\kappa} - 1\right) s.$$

Earlier in the horizon, the option value of waiting raises thresholds above their terminal value, because rejecting preserves the opportunity to draw a better partner in future subperiods.

PA sits between the two benchmark cases. Relative to unawareness, agents now correctly anticipate that partners evaluate themselves through a distorted lens $s_l = \lambda_l t_l$, but they continue misperceive their own market value. This shifts the *perceived* mutual acceptance interval by the factor $1/\lambda_l$ and breaks the single symmetric recursion unless $\lambda_A = \lambda_B$. Relative to FA, the strategic fixed point is still defined through the opponent's cutoffs, but the beliefs feeding into continuation values are systematically misaligned with the true mutual-acceptance event.

4 Cheap Search

We now consider the consequences of cheap search. Cheap search is modeled by letting the number of sub-periods per unit time grow. Formally, we write the total horizon as $N \equiv nm$ and let $m \to \infty$, while holding the macro-horizon n fixed. The per-subperiod discount factor is $\hat{\delta} \equiv \delta^{1/m}$, so that overall impatience over a macro-period remains fixed at δ. We will focus on subperiods with a large remaining horizon ($H_i \geq \eta N$), i.e. away from the terminal layer where end effects dominate.

The following lemma shows that when search is sufficiently cheap, perceived cutoffs cannot lie far below the relevant upper feasibility frontier. However, they might lie strictly above it. The lemma applies to all three cases: UA, FA, and PA, and it holds for any $\lambda_A, \lambda_B > 0$.

Lemma 2 (Zero measure perceived acceptance) *In all three cases, UA, FA, and PA, perceived mutual acceptance has zero measure on compacts away from the terminal layer in the limit for $m \to \infty$.*

Formally, fix $\eta \in (0,1)$ and $0 < \underline{s} < \bar{s} < \infty$, and write $t_k = s/\lambda_k$. Define $\varepsilon_m \equiv m^{-1/2}$. There exists a constant $C < \infty$ such that for all m sufficiently large, all i with $H_i \geq \eta N$, and uniformly over $s \in [\underline{s}, \bar{s}]$:

1. ***UA:*** $\left(\rho_{l,i}^{-1}[s] - \rho_{k,i}[s]\right)_+ \leq C\,\varepsilon_m$
2. ***FA:*** $\left(\frac{1}{\lambda_l}\rho_{l,i}^{-1}[t_k] - \rho_{k,i}[s]\right)_+ \leq C\,\varepsilon_m$
3. ***PA:*** $\left(\frac{1}{\lambda_l}\rho_{l,i}^{-1}[s] - \rho_{k,i}[s]\right)_+ \leq C\,\varepsilon_m$.

Proof. Let $l \neq k$ and abbreviate

$$x \equiv \rho_{k,i}[s], \qquad y \equiv \frac{1}{\lambda_l}\rho_{l,i}^{-1}[s].$$

If $(y-x)_+ = 0$ we are done, so suppose $x < y$.

Partial Awareness:

Step 1 (bounds away from 0 on the relevant region). Let $a \equiv (2-\kappa)/\kappa > 0$. Matching with partner type t yields flow payoff $\frac{\kappa}{2}(s+t)$, which is strictly below s whenever $t < as$. Hence accepting $t < as$ is strictly dominated, so

$$x = \rho_{k,i}[s] \geq as \geq a\underline{s}. \tag{6}$$

Since $x < y$ by assumption, we have $y > a\underline{s}$. Applying (6) to $\rho_{l,i}$ gives $\rho_{l,i}[u] \geq au$, hence $\rho_{l,i}^{-1}[s] \leq s/a \leq \bar{s}/a$, so

$$y = \frac{1}{\lambda_l}\rho_{l,i}^{-1}[s] \leq \frac{\bar{s}}{a\lambda_l}. \tag{7}$$

Thus y lies in a compact interval $[a\underline{s}, \bar{s}/(a\lambda_l)] \subset (0, \infty)$.

Step 2 (a feasible "wait for the top band" strategy). Recall $\varepsilon_m \equiv m^{-1/2}$. Consider the following feasible deviation at state (i, s): reject all draws until the opponent type lies in $[y - \varepsilon_m, y]$, then accept. This is feasible from state (i, s) onward: by definition of y, any $t_l \leq y$ is believed to accept k at perceived type s, and by Lemma 12 in the Appendix, the opponent's cutoff is weakly decreasing over time, so such types remain believed to accept in all subsequent subperiods.

Let T be the number of draws until $[y - \varepsilon_m, y]$ is first hit. Then T is geometric with success probability

$$p_m = F[y] - F[y - \varepsilon_m] = \int_{y-\varepsilon_m}^{y} f[t]\, dt.$$

By Step 1, for large m we have $y - \varepsilon_m \geq \frac{1}{2}a\underline{s}$. Hence by continuity and strict positivity of f, there exists $\underline{f} > 0$ (depending only on $\underline{s}, \bar{s}$ and primitives) such that $f[t] \geq \underline{f}$ on $[\frac{1}{2}a\underline{s}, \bar{s}/(a\lambda_l)]$, so

$$p_m \geq \underline{f}\, \varepsilon_m. \tag{8}$$

Write $\hat{\delta} \equiv \delta^{1/m}$. For a geometric T,

$$\mathbb{E}\left[\hat{\delta}^{T-1}\right] = \sum_{t\geq 1}(1-p_m)^{t-1}p_m\hat{\delta}^{t-1} = \frac{p_m}{1-(1-p_m)\hat{\delta}} = \frac{p_m}{(1-\hat{\delta})+p_m\hat{\delta}}.$$

Therefore

$$1-\mathbb{E}\Big[\hat{\delta}^{T-1}\Big]=\frac{1-\hat{\delta}}{(1-\hat{\delta})+p_m\hat{\delta}}\le\frac{1-\hat{\delta}}{p_m\hat{\delta}}\le\frac{-\ln\delta}{\delta}\cdot\frac{1}{m\,p_m}\le C\cdot\frac{1}{m\,\varepsilon_m},$$

where we used $1-\delta^{1/m}\le(-\ln\delta)/m$, $\hat{\delta}\ge\delta$, and (8). Moreover,

$$\mathbb{E}\Big[\hat{\delta}^{T-1}(T-1)\Big]=\sum_{t\ge1}(t-1)(1-p_m)^{t-1}p_m\hat{\delta}^{t-1}=\frac{(1-p_m)\hat{\delta}\,p_m}{\big(1-(1-p_m)\hat{\delta}\big)^2}\le\frac{1}{\delta^2}\cdot\frac{1}{p_m}\le C\cdot\frac{1}{\varepsilon_m}.$$

Step 3 (compare the cutoff with the deviation). At the cutoff, the agent is indifferent: rejecting at (i,s) yields the equilibrium continuation value, and that continuation value is at least the payoff of the feasible deviation in Step 2. Hence the equilibrium value from *accepting* a partner of type x at (i,s) is at least the deviation value.

Accepting at (i,s) with partner type x yields (since matching is absorbing)

$$\alpha_i\cdot\frac{\kappa}{2}(s+x).$$

For the deviation, ignore (positive) single payoffs while waiting. Conditional on T, acceptance occurs at some $t_l\ge y-\varepsilon_m$ after $T-1$ draws, so the deviation payoff is at least

$$\mathbb{E}\Big[\hat{\delta}^{T-1}\alpha_{i+T-1}\Big]\cdot\frac{\kappa}{2}(s+y-\varepsilon_m),$$

where α_{i+T-1} is the remaining annuity factor at the acceptance date. Since each term in α_i is at most 1, we have $\alpha_{i+T-1}\ge\alpha_i-(T-1)$, hence

$$\mathbb{E}\Big[\hat{\delta}^{T-1}\alpha_{i+T-1}\Big]\ge\alpha_i\mathbb{E}\Big[\hat{\delta}^{T-1}\Big]-\mathbb{E}\Big[\hat{\delta}^{T-1}(T-1)\Big].$$

When $H_i\ge\eta N$, we have $\alpha_i\ge c\,m$ for a constant $c\equiv(1-\delta^{\eta n})/(-\ln\delta)>0$. Combining this with the bounds from Step 2 gives, for large m,

$$\frac{1}{\alpha_i}\mathbb{E}\Big[\hat{\delta}^{T-1}\alpha_{i+T-1}\Big]\ge1-C\cdot\frac{1}{m\,\varepsilon_m}.$$

Therefore

$$s+x\ \ge\ \left(1-C\cdot\frac{1}{m\,\varepsilon_m}\right)(s+y-\varepsilon_m),$$

which implies

$$0\le y-x\le\varepsilon_m+C\cdot\frac{s+y}{m\,\varepsilon_m}\le C'\left(\varepsilon_m+\frac{1}{m\,\varepsilon_m}\right).$$

Since $\varepsilon_m=m^{-1/2}$, the right-hand side is $O(m^{-1/2})$, uniformly in i (as long as $H_i\ge\eta N$) and uniformly over $s\in[\underline{s},\bar{s}]$. This proves the partial-awareness bound.

Unawareness:

Redo steps 1–3 with (perceived) $\lambda_l = 1$ (so $y \equiv \rho_{l,i}^{-1}[s]$). This gives the unawareness bound.

Full Awareness:

Redo steps 1–3 with $t_k \equiv s/\lambda_k$ and $y \equiv (1/\lambda_l)\rho_{l,i}^{-1}[t_k]$. In Step 1, the same dominated-acceptance bound yields $x \geq as$ while $\rho_{l,i}^{-1}[t_k] \leq t_k/a$, so $y \leq \bar{s}/(a\lambda_k\lambda_l)$. Steps 2–3 are unchanged, and the same conclusion follows.

This establishes the lemma. ■

4.1 Unawareness and the Paradox of Plenty

In UA, an agent with self-perception s in subperiod i is (subjectively) accepted by a type-t counter part iff $t \leq \rho_i^{-1}[s]$ (equivalently $\rho_i[t] \leq s$). Therefore, any best reply that sets one's cutoff strictly above $\rho_i^{-1}[s]$ cannot create any match (since all $t > \rho_i^{-1}[s]$ reject). Moreover, replacing such a cutoff by $\rho_i^{-1}[s]$ can only add matches at the boundary $t = \rho_i^{-1}[s]$, which has probability zero since F admits a continuous density. More generally, whenever multiple best replies differ only on a probability-zero boundary, we adopt the smallest cutoff as a canonical tie-break. Thus, we may (and henceforth do) adopt the canonical best reply in which

$$\rho_i[s] \leq \rho_i^{-1}[s] \qquad \forall s > 0. \tag{9}$$

Consider the UA SSPE of Proposition 1.1. When search is cheap, rejecting and drawing again has negligible value cost relative to the gain from a slightly better partner. Hence any slack in the feasible acceptance interval is eliminated, and the cutoff tracks one's perceived type. As a result, the perceived mutual-acceptance interval $[\rho_i[s], \rho_i^{-1}[s]]$ shrinks toward the knife-edge condition $t \approx s$ on any compact region away from 0.

Lemma 3 *Suppose agents are unaware. Away from the terminal layer, the perceived mutual acceptance region becomes knife-edge as $m \to \infty$, collapsing uniformly to the 45-degree line on compacts.*

Formally, for any $0 < \underline{s} < \bar{s} < \infty$ and $\eta \in (0,1)$ there exists a sequence $\varepsilon_m \downarrow 0$ such that for all sufficiently large m, for all subperiods i with $H_i \geq \eta N$, and for all $s \in [\underline{s}, \bar{s}]$,

$$(1-\varepsilon_m)s \leq \rho_i[s] \leq (1+\varepsilon_m)s, \qquad (1-\varepsilon_m)s \leq \rho_i^{-1}[s] \leq (1+\varepsilon_m)s. \tag{10}$$

Proof. By the canonical tie-breaking (9) and the strict monotonicity of $\rho_i[\cdot]$ (Lemma 11), we have

$$\rho_i[s] \leq s \leq \rho_i^{-1}[s] \qquad \forall s > 0. \tag{11}$$

Indeed, if $\rho_i[s] > s$, then strict monotonicity implies $\rho_i^{-1}[s] \leq s$, hence $\rho_i[s] > \rho_i^{-1}[s]$, contradicting (9).

Under unawareness the equilibrium is symmetric, so both sides use the same cutoff $\rho_i[\cdot]$ (Proposition 1.1). The compact-width bound from (the proof of) Lemma 2 implies that there exist $C < \infty$ and m_0 (depending on the compact and primitives) such that for all $m \geq m_0$,

$$\sup_{s \in [\underline{s}, \bar{s}]} \left(\rho_i^{-1}[s] - \rho_i[s]\right) \leq C\, m^{-1/4}.$$

Combining this with (11) yields, for all $m \geq m_0$ and all $s \in [\underline{s}, \bar{s}]$,

$$|\rho_i[s] - s| \leq C\, m^{-1/4}, \qquad |\rho_i^{-1}[s] - s| \leq C\, m^{-1/4}.$$

Setting $\varepsilon_m \equiv (C/\underline{s})\, m^{-1/4}$ (so $\varepsilon_m \downarrow 0$) and using $s \geq \underline{s}$ on the compact gives (10). ■

Lemma 3 implies that, away from the terminal layer, we may approximate ρ_i and ρ_i^{-1} by the identity map. Plugging this approximation into the expression for M_i yields the simple strip $t_B \gtrsim \lambda_A t_A$ and $t_B \lesssim \frac{1}{\lambda_B} t_A$. When $\lambda_A \lambda_B > 1$, the lower bound eventually exceeds the upper bound on any compact bounded away from 0, so realized mutual acceptance collapses. When $\lambda_A \lambda_B < 1$, the strip has nonzero width and mutual acceptance remains possible.

Proposition 4 (Paradox of Plenty: Unawareness) *When search is sufficiently cheap, i.e., m is large enough, the actual mutual acceptance region away from the terminal layer is empty on any compact bounded away from 0 iff $\lambda_A \lambda_B > 1$.*

Formally, for any $0 < \underline{t} < \bar{t} < \infty$ and $\eta \in (0, 1)$:

1. *If $\lambda_A \lambda_B > 1$, there exists an m_0 such that for all $m \geq m_0$ and all i with $H_i \geq \eta N$,*

$$M_i \cap ([\underline{t}, \bar{t}] \times [\underline{t}, \bar{t}]) = \varnothing.$$

2. *If $\lambda_A \lambda_B < 1$, then for all m and all i, $M_i \neq \varnothing$. Specifically, for any $t_A \in [\underline{t}, \bar{t}]$ and $t_B \in \left[\lambda_A t_A, \frac{1}{\lambda_B} t_A\right]$, $(t_A, t_B) \in M_i$.*

Proof. Fix $0 < \underline{t} < \bar{t} < \infty$ and choose ε_m as in Lemma 3 on the compact $[\underline{s}, \bar{s}]$, where $\underline{s} \equiv \min\{\underline{t}, \lambda_A \underline{t}\}$ and $\bar{s} \equiv \max\{\bar{t}, \lambda_A \bar{t}\}$. Then (10) holds for all sufficiently large m and all i with $H_i \geq \eta N$.

(i) $\lambda_A \lambda_B > 1$**.** Fix $t_A \in [\underline{t}, \bar{t}]$. The lower bound in M_i satisfies $\rho_i[\lambda_A t_A] \geq (1 - \varepsilon_m)\lambda_A t_A$, while the upper bound satisfies $\frac{1}{\lambda_B}\rho_i^{-1}[t_A] \leq (1 + \varepsilon_m)\frac{1}{\lambda_B} t_A$. Thus a necessary condition for M_i to contain any point with this t_A is

$$(1 - \varepsilon_m)\lambda_A t_A \leq (1 + \varepsilon_m)\tfrac{1}{\lambda_B} t_A \quad \iff \quad \lambda_A \lambda_B \leq \frac{1 + \varepsilon_m}{1 - \varepsilon_m}.$$

If $\lambda_A \lambda_B > 1$, choose m so large that $\frac{1+\varepsilon_m}{1-\varepsilon_m} < \lambda_A \lambda_B$. Then the inequality fails for every $t_A > 0$, so $M_i \cap ([\underline{t}, \bar{t}] \times [\underline{t}, \bar{t}]) = \varnothing$.

(ii) $\lambda_A \lambda_B < 1$**.** Under the canonical tie-breaking (9), we have (11), i.e. $\rho_i[x] \leq x \leq \rho_i^{-1}[x]$ for all $x > 0$. Fix any m, any i, and any $t_A \in [\underline{t}, \bar{t}]$. For any $t_B \in [\lambda_A t_A,\ \frac{1}{\lambda_B} t_A]$,

$$\rho_i[\lambda_A t_A] \leq \lambda_A t_A \leq t_B \qquad \text{and} \qquad t_B \leq \tfrac{1}{\lambda_B} t_A \leq \tfrac{1}{\lambda_B}\rho_i^{-1}[t_A],$$

so $(t_A, t_B) \in M_i$. This proves (ii). ■

The proposition derives a stark, multiplicative condition for breakdown under UA. In the frictionless-search regime, each side's self-perception scales up the type it thinks it has, and this scaling enters the two one-sided acceptance constraints in opposite directions. If the product $\lambda_A \lambda_B$ exceeds one, these constraints become incompatible away from very small types and away from the end of the horizon, so the matching set becomes empty.

The perceived set of mutually acceptable partners for t_k in subperiod i is

$$\mathcal{M}_{k,i}^{\text{perc}}[t_k] \equiv \left\{ t_l \geq 0 : \rho_i[\lambda_k t_k] \leq t_l \leq \rho_i^{-1}[\lambda_k t_k] \right\},$$

while the actual set of mutually acceptable partners is

$$\mathcal{M}_{k,i}^{\text{act}}[t_k] \equiv \left\{ t_l \geq 0 : \rho_i[\lambda_k t_k] \leq t_l \leq \tfrac{1}{\lambda_l}\rho_i^{-1}[\lambda_k t_k] \right\}.$$

We say that t_k is a *false hoper* in subperiod i if $\mathcal{M}_{k,i}^{\text{perc}}[t_k] \neq \varnothing$ but $\mathcal{M}_{k,i}^{\text{act}}[t_k] = \varnothing$. That is, false hopers continue to search because they (mistakenly) believe they have a chance of matching, although in truth they do not.

Corollary 5 (False Hopers) *If $\lambda_A \lambda_B > 1$, then cheap search turns almost everyone into a false hoper.*

Formally, if $\lambda_A\lambda_B > 1$ then, for every $t_k > 0$, there exists an m_0 such that for all $m \geq m_0$ and every i with $H_i \geq \eta N$,

$$\mathcal{M}^{\text{perc}}_{k,i}[t_k] \neq \varnothing \qquad \textit{and} \qquad \mathcal{M}^{\text{act}}_{k,i}[t_k] = \varnothing \ .$$

Proof. By the canonical tie-breaking (9), we have $\rho_i[s] \leq \rho_i^{-1}[s]$ for all $s > 0$, hence $\mathcal{M}^{\text{perc}}_{k,i}[t_k] \neq \varnothing$ for every $t_k > 0$ and every i. Under $\lambda_A\lambda_B > 1$, Proposition 4(i) shows that for m large enough and all interior layers i with $H_i \geq \eta N$, the actual mutual acceptance constraints become incompatible on $[\underline{t}, \bar{t}]$, i.e. $\mathcal{M}^{\text{act}}_{k,i}(t_k) = \varnothing$ for all $t_k \in [\underline{t}, \bar{t}]$ and $k \in \{A, B\}$. This implies the pointwise false-hoper claim in the corollary. ■

4.2 Cheap Search with Full Awareness

We now state the full-FA analogue of the cheap-search geometry result: for large m and away from the terminal layer, mutual acceptance either becomes knife-edge or becomes empty, depending on κ and the overconfidence wedges λ_A, λ_B.

Proposition 6 (Cheap-search Matching: Awareness) *If $\lambda_A\lambda_B \leq 1/\left(\frac{2}{\kappa} - 1\right)^2$ then, away from the terminal layer, the mutual acceptance region is knife-edge or empty for $m \to \infty$.*

If $\lambda_A\lambda_B > 1/\left(\frac{2}{\kappa} - 1\right)^2$, then the mutual acceptance region is empty on $(0, \infty)^2$, for all $m \in \mathbb{N}$ and all $i \leq N$.

Formally, for any $0 < \underline{t} < \bar{t} < \infty$ and any $\eta \in (0, 1)$ there exists a sequence $\varepsilon_m \downarrow 0$ such that for all sufficiently large m and all i with $H_i \geq \eta N$:

1. *If $\lambda_A\lambda_B \leq 1/\left(\frac{2}{\kappa} - 1\right)^2$, then*

$$M_i \cap [\underline{t}, \bar{t}]^2 \subseteq \left\{(t_A, t_B) \in [\underline{t}, \bar{t}]^2 : t_B \in \left[\rho_{A,i}[\lambda_A t_A],\ \rho_{A,i}[\lambda_A t_A] + \varepsilon_m\right]\right\}.$$

2. *If $1/\left(\frac{2}{\kappa} - 1\right)^2 < \lambda_A\lambda_B$, then*

$$M_i \cap [\underline{t}, \bar{t}]^2 = \varnothing.$$

Proof. A useful bound. For any side k and any subperiod $i \leq N$, the cutoff satisfies

$$\rho_{k,i}[s] \geq \left(\frac{2}{\kappa} - 1\right) s \qquad \forall s > 0. \tag{12}$$

To see this, use the indifference condition defining $\rho_{k,i}[s]$:

$$\alpha_i \frac{\kappa}{2}\left(s + \rho_{k,i}[s]\right) = s + \hat{\delta} V_{k,i+1}[s].$$

Even if the agent were to remain single forever from $i+1$ onward, he would obtain $\alpha_{i+1}s$, hence $V_{k,i+1}[s] \geq \alpha_{i+1}s$. Therefore $s + \hat{\delta}V_{k,i+1}[s] \geq s + \hat{\delta}\alpha_{i+1}s = \alpha_i s$, and dividing by α_i yields $\frac{\kappa}{2}(s + \rho_{k,i}[s]) \geq s$, i.e. (12).

(1) $\boldsymbol{\lambda_A \lambda_B \leq 1/\left(\frac{2}{\kappa} - 1\right)^2}$. Consider any subperiod i with $H_i \geq \eta N$, and for $t_A \in [\underline{t}, \bar{t}]$, let

$$x \equiv \rho_{A,i}[\lambda_A t_A], \qquad y \equiv \frac{1}{\lambda_B}\rho_{B,i}^{-1}[t_A].$$

By Lemma 2 (Full FA), applied on the compact $[\lambda_A \underline{t}, \lambda_A \bar{t}]$, there exists a sequence $\varepsilon_m \downarrow 0$ such that for all sufficiently large m,

$$(y - x)_+ \leq \varepsilon_m$$

uniformly over $t_A \in [\underline{t}, \bar{t}]$ and all such i. Therefore, if $(t_A, t_B) \in M_i \cap [\underline{t}, \bar{t}]^2$ then necessarily $x \leq t_B \leq y \leq x + \varepsilon_m$, which yields the stated inclusion.

(2) $\boldsymbol{\lambda_A \lambda_B > 1/\left(\frac{2}{\kappa} - 1\right)^2}$. Take any (t_A, t_B) with $t_A > 0$ and suppose it were mutually acceptable in subperiod i. By (12),

$$t_B \geq \rho_{A,i}[\lambda_A t_A] \geq \left(\frac{2}{\kappa} - 1\right)\lambda_A t_A \quad \text{and} \quad t_A \geq \rho_{B,i}[\lambda_B t_B] \geq \left(\frac{2}{\kappa} - 1\right)\lambda_B t_B.$$

Combining the two inequalities gives $t_B \geq \left(\frac{2}{\kappa} - 1\right)^2 \lambda_A \lambda_B t_B$, which is impossible for $t_B > 0$ when $\lambda_A \lambda_B > 1/\left(\frac{2}{\kappa} - 1\right)^2$. Hence $M_i \cap [\underline{t}, \bar{t}]^2 = \varnothing$ for every i; in particular, the stated emptiness holds for all sufficiently large m and all i with $H_i \geq \eta N$. ■

Part (1) shows that, whenever $\lambda_A \lambda_B \leq 1/\left(\frac{2}{\kappa} - 1\right)^2$, any mutually acceptable pair in an interior subperiod must lie within an $o(1)$ band of the (endogenous) mutual-feasibility frontier $t_B = \rho_{A,i}[\lambda_A t_A]$ (equivalently $t_B = \frac{1}{\lambda_B}\rho_{B,i}^{-1}[t_A]$). Part (2) identifies the parameter region in which these bounds are incompatible for any $t_B > 0$, so that mutual acceptance is empty (in particular, on any compact bounded away from 0).

4.3 Partial Awareness and the Paradox of Plenty

Fix $0 < \underline{t} < \bar{t} < \infty$ and $\eta \in (0, 1)$. In Partial Awareness (PA), each agent k best-responds to a perceived state $s_k = \lambda_k t_k$ while meeting opportunities are indexed by market types t_k. This creates a wedge between *perceived* and *actual* feasibility.

To illustrate, conditional on t_A (so $s_A = \lambda_A t_A$), agent A's *perceived* feasibility interval for t_B is

$$\mathcal{M}^{\text{perc}}_{A,i}[t_A] \equiv \left[\rho_{A,i}[s_A],\ \frac{1}{\lambda_B}\rho^{-1}_{B,i}[s_A]\right],$$

whereas the *actual* feasibility interval is

$$\mathcal{M}^{\text{act}}_{A,i}[t_A] \equiv \left[\rho_{A,i}[s_A],\ \frac{1}{\lambda_B}\rho^{-1}_{B,i}[t_A]\right].$$

(Analogously, conditional on t_B with $s_B = \lambda_B t_B$, the perceived and actual intervals for A are obtained by swapping A and B.) Notice that in PA, actual feasibility implies perceived feasibility pointwise ($\mathcal{M}^{\text{act}}_{A,i}[t_A] \subset \mathcal{M}^{\text{perc}}_{A,i}[t_A]$).

Lemma 2.3 shows that, away from the terminal layer and on compacts, the *positive part* of the perceived width becomes knife-edge as $m \to \infty$. Actual feasibility may nevertheless vanish because the actual upper endpoint evaluates $\rho^{-1}_{l,i}[\cdot]$ at market type t_k rather than perceived type $s_k = \lambda_k t_k$.

To turn knife-edge *perceived* slack into *empty actual feasibility* on a fixed compact, we would need an order-one separation in the opponent inverse cutoff between the market type t and the perceived type $\lambda_k t$. Otherwise the s–to–t correction in the actual upper endpoint could shrink along with the perceived slack.

Fix distinct $k, l \in \{A, B\}$ with $k \neq l$. We record this requirement as the following increment condition.

Condition 7 (Increment on compacts) *Scaling t by λ_k shifts the opponent inverse cutoff by an order-one amount.*

(INC): *Formally, there exists an $\varepsilon > 0$ such that for all m sufficiently large and all i with $H_i \geq \eta N$,*

$$\rho^{-1}_{l,i}[\lambda_k t] - \rho^{-1}_{l,i}[t] \geq \varepsilon \qquad \forall\, t \in [\underline{t}, \bar{t}].$$

A more transparent ρ–form. Since $\rho_{l,i}$ is increasing, (INC) is equivalent to a one-sided growth restriction on the cutoff itself (no inverses). Let $S_{l,i} \equiv \rho^{-1}_{l,i}([\underline{t}, \bar{t}])$. Then (INC) holds with constant ε if and only if, for every $s \in S_{l,i}$ such that $s + \varepsilon$ lies in the domain of $\rho_{l,i}$,

$$\rho_{l,i}[s + \varepsilon] \leq \lambda_k\, \rho_{l,i}[s].$$

In words: shifting the opponent *inverse* cutoff by ε is the same as requiring that the cutoff map $\rho_{l,i}$ cannot multiply by the factor λ_k over an s–increment of length ε on the relevant band.

A simple sufficient condition (and how it can fail). If $\rho_{l,i}$ is absolutely continuous on $\rho_{l,i}^{-1}([\underline{t}, \lambda_k \bar{t}])$ and $\rho'_{l,i}(s) \leq L$ a.e. there, then taking $\varepsilon \equiv (\lambda_k - 1)\underline{t}/L$ implies $\rho_{l,i}[s + \varepsilon] \leq \rho_{l,i}[s] + L\varepsilon \leq \lambda_k \rho_{l,i}[s]$ for all $s \in S_{l,i}$, hence (INC). Conversely, (INC) is violated whenever $\rho_{l,i}$ exhibits a sufficiently steep local segment on the band, i.e., whenever one can find s with $\rho_{l,i}[s + \delta] > \lambda_k \rho_{l,i}[s]$ for some small δ; Proposition 10 shows that such steepening is generic under PA.

Condition 7 pins down the *size* of the PA misperception correction on the compact: replacing t_k by $s_k = \lambda_k t_k$ shifts the opponent inverse cutoff by at least ε. Since k's actual upper endpoint evaluates $\rho_{l,i}^{-1}[\cdot]$ at t_k whereas k's perceived upper endpoint evaluates it at s_k, the actual signed slack equals the perceived signed slack minus $(1/\lambda_l)\big(\rho_{l,i}^{-1}[s_k] - \rho_{l,i}^{-1}[t_k]\big) \leq -\varepsilon/\lambda_l$. Thus, once perceived slack is small, the actual feasibility interval must be empty.

Proposition 8 (Paradox of Plenty under Partial Awareness) *Under INC,*

$$M_i \cap [\underline{t}, \bar{t}]^2 = \varnothing.$$

for all m sufficiently large and all i with $H_i \geq \eta N$.

Proof. Since $s_k = \lambda_k t_k$ ranges over the compact $[\lambda_k \underline{t}, \lambda_k \bar{t}]$, Lemma 2.3 implies that there exists an m_0 such that for all $m \geq m_0$ and all i with $H_i \geq \eta N$,

$$\left(\frac{1}{\lambda_l}\rho_{l,i}^{-1}[s_k] - \rho_{k,i}[s_k]\right)_+ \leq \frac{\varepsilon}{2\lambda_l} \qquad \forall\, t_k \in [\underline{t}, \bar{t}].$$

Under (INC), for the same m and i we also have $\rho_{l,i}^{-1}[s_k] - \rho_{l,i}^{-1}[t_k] \geq \varepsilon$ for all $t_k \in [\underline{t}, \bar{t}]$. Therefore the signed width of k's *actual* feasibility interval satisfies

$$\frac{1}{\lambda_l}\rho_{l,i}^{-1}[t_k] - \rho_{k,i}[s_k] = \left(\frac{1}{\lambda_l}\rho_{l,i}^{-1}[s_k] - \rho_{k,i}[s_k]\right) - \frac{1}{\lambda_l}\left(\rho_{l,i}^{-1}[s_k] - \rho_{l,i}^{-1}[t_k]\right) \leq \frac{\varepsilon}{2\lambda_l} - \frac{\varepsilon}{\lambda_l} < 0,$$

so k's actual feasibility interval is empty for each t_k in the compact. Finally, if both one-sided actual feasibility intervals are empty on $[\underline{t}, \bar{t}]$ in subperiod i, then no pair in $[\underline{t}, \bar{t}]^2$ can satisfy mutual acceptance, so $M_i \cap [\underline{t}, \bar{t}]^2 = \varnothing$. ■

Interpretation. In PA, perceived feasibility on compacts becomes knife-edge as $m \to \infty$ (Lemma 2.3). If, in addition, the misperception correction in the actual upper endpoint remains uniformly bounded away from 0 (INC), then this correction dominates the vanishing perceived slack, so no pair in $[\underline{t}, \bar{t}]^2$ can satisfy mutual acceptance on interior layers.

The Paradox-of-Plenty conclusion on compacts hinges on a single extra ingredient beyond knife-edge perceived feasibility: Condition 7 (INC), which enforces an *order-one* separation of inverse cutoffs between t and $\lambda_k t$.

Rather than treating (INC) as innocuous, we show that in PA it is *not robust*: arbitrarily small perturbations of primitives can generate regions in which the cutoff becomes locally steep, collapsing inverse increments on the compact. The next lemma gives a simple operational criterion for failure of (INC); Proposition 10 then upgrades this to an open-and-dense (generic) failure result.

Lemma 9 (Steep cutoffs kill (INC)) *If the cutoff is locally sufficiently steep at some point mapping into $[\underline{t}, \bar{t}]$, then the corresponding inverse increment is $< \varepsilon$.*

Fix $\varepsilon > 0$. Let $\rho_{l,i}[\cdot]$ be strictly increasing and absolutely continuous on an interval containing $\rho_{l,i}^{-1}([\underline{t}, \lambda_k \bar{t}])$. If there exists a point s_0 at which $\rho_{l,i}$ is differentiable, with $t_0 \equiv \rho_{l,i}[s_0] \in [\underline{t}, \bar{t}]$ and

$$\rho'_{l,i}[s_0] > \frac{2(\lambda_k - 1)\bar{t}}{\varepsilon},$$

then $\rho_{l,i}^{-1}[\lambda_k t_0] - \rho_{l,i}^{-1}[t_0] < \varepsilon$.

Proof. By differentiability at s_0, there exists $\delta_0 > 0$ such that for all $h \in (0, \delta_0)$,

$$\rho_{l,i}[s_0 + h] - \rho_{l,i}[s_0] > \tfrac{1}{2}\, \rho'_{l,i}[s_0]\, h.$$

Choose $h \equiv 2(\lambda_k - 1)t_0 / \rho'_{l,i}[s_0]$. Since $t_0 \leq \bar{t}$ and $\rho'_{l,i}[s_0] > 2(\lambda_k - 1)\bar{t}/\varepsilon$, we have $h < \varepsilon$; taking $h < \delta_0$ as well, the inequality gives

$$\rho_{l,i}[s_0 + h] > t_0 + (\lambda_k - 1)t_0 \;=\; \lambda_k t_0.$$

Because $\rho_{l,i}$ is strictly increasing, this implies $\rho_{l,i}^{-1}[\lambda_k t_0] < s_0 + h$ and $\rho_{l,i}^{-1}[t_0] = s_0$, hence $\rho_{l,i}^{-1}[\lambda_k t_0] - \rho_{l,i}^{-1}[t_0] < h < \varepsilon$. ■

We next show that local steepness (and hence failure of (INC)) can be induced by an arbitrarily small perturbation of primitives. The argument exploits the fact that on interior layers the effective horizon α_i is large, so boundary terms in $V'_{l,i}$ can be amplified.

Primitive space and genericity. Consider the PA SSPE of Proposition 1. Fix distinct $k, l \in \{A, B\}$ with $\lambda_k > 1$, and set $\bar{\lambda} \equiv \max\{\lambda_A, \lambda_B\}$. Let $\mathcal{P}$ be the set of primitives

$p \equiv (f, \kappa, \hat{\delta}, \lambda_A, \lambda_B)$ with $f \in C^1$ on $[0, \bar{\lambda}\bar{t}]$ strictly positive and bounded, endowed with the product topology generated by the C^1 norm on f and the Euclidean norm on parameters. For $r \in \mathbb{N}$ define

$$\mathcal{O}_r \equiv \bigcup_{m \geq r} \left\{ p \in \mathcal{P} : \exists\, i \text{ with } H_i \geq \eta N,\ \exists\, t \in [\underline{t}, \bar{t}] \text{ s.t. } \rho_{l,i}^{-1}[\lambda_k t] - \rho_{l,i}^{-1}[t] < \frac{1}{r} \right\},$$

and let $\mathcal{O} \equiv \bigcap_{r \geq 1} \mathcal{O}_r$. Thus $p \in \mathcal{O}_r$ means that (INC) fails at scale $1/r$ on some interior layer for some sufficiently large m.

Standing continuity assumption. The SSPE correspondence $p \mapsto (\rho_{l,i}^{-1}, \rho'_{l,i})$ is continuous (on compacts, for fixed m, i) at every $p \in \mathcal{P}$. This is a standard consequence of uniqueness of the fixed point and interior differentiability of best responses; we only use it to pass strict inequalities to nearby primitives.

Proposition 10 (Generic failure of (INC) in PA) *(INC) is non-robust in PA: for every accuracy level $1/r$, arbitrarily small perturbations of primitives lead (for some sufficiently large m) to an interior layer and a point $t \in [\underline{t}, \bar{t}]$ at which the increment falls below $1/r$.*

Formally, for each $r \in \mathbb{N}$, the set $\mathcal{O}_r$ is open and dense in $\mathcal{P}$. Hence the residual set $\mathcal{O}$ is comeager in $\mathcal{P}$, and every $p \in \mathcal{O}$ violates Condition 7 (INC).

Proof. *Openness.* Fix r and suppose $p \in \mathcal{O}_r$. Then there exist $m \geq r$, an interior layer i with $H_i \geq \eta N$, and a $t \in [\underline{t}, \bar{t}]$ such that $\rho_{l,i}^{-1}[\lambda_k t] - \rho_{l,i}^{-1}[t] < 1/r$. Since the inequality is strict and the map $p \mapsto \rho_{l,i}^{-1}[\cdot]$ is continuous on the compact $[\underline{t}, \lambda_k \bar{t}]$, the same strict inequality holds for all primitives in a sufficiently small neighborhood of p. Therefore $\mathcal{O}_r$ is open.

Density (explicit perturbation). Fix $p \in \mathcal{P}$ and a neighborhood $\mathcal{U}$ of p. Choose m so large that there exists an interior layer i with $H_i \geq \eta N$ and α_i arbitrarily large.[1] Pick a point $s_0 \in [\underline{t}, \bar{\lambda}\bar{t}]$ with $U_{l,i}(s_0) > \rho_{l,i}[s_0]$ (so $\pi_{l,i}(s_0) > 0$).[2]

Let $u_0 \equiv U_{l,i}(s_0)$ and choose a small interval $I \subset (\underline{t}, \bar{\lambda}\bar{t})$ with $u_0 \in I \subset (\rho_{l,i}[s_0], U_{l,i}(s_0))$. Let ψ be a nonnegative C^1 bump supported on I with $\int \psi = 1$, and let χ be a nonnegative C^1 bump supported outside the acceptance band of (l, i) on $[\underline{t}, \bar{\lambda}\bar{t}]$ with $\int \chi = 1$. For $\theta > 0$ define the perturbed density

$$f^{(\theta)}(t) \ \equiv\ f(t) \ +\ \theta\big(\psi(t) - \chi(t)\big).$$

[1] Recall $\alpha_i \equiv \sum_{h=0}^{H_i - 1} \hat{\delta}^h$ and, with $\hat{\delta} = \delta^{1/m}$, we have $\alpha_i \to \infty$ uniformly across interior layers as $m \to \infty$.

[2] If $\pi_{l,i} \equiv 0$ on $[\underline{t}, \bar{\lambda}\bar{t}]$ for every interior i, then matching is already empty on compacts and the paradox mechanism is vacuous.

For θ small, $f^{(\theta)}$ remains strictly positive and C^1 and hence $p^{(\theta)} \equiv (f^{(\theta)}, \kappa, \hat{\delta}, \lambda_A, \lambda_B) \in \mathcal{U}$.

Because the perturbation shifts probability mass *into* the acceptance band around u_0 while keeping total mass unchanged, it strictly increases the match probability $\pi_{l,i}(s_0)$ at first order in θ. Differentiating the bootstrap value representation (5) w.r.t. s and applying the Leibniz rule shows that $V'_{l,i}[s_0]$ contains a boundary term proportional to $\alpha_i \kappa \pi'_{l,i}(s_0)$, where $\pi'_{l,i}(s_0)$ depends linearly on f through the endpoint evaluations $f(U_{l,i}(s_0))$ and $f(\rho_{l,i}[s_0])$. Under the present bump perturbation, this boundary term changes by an amount of order $\alpha_i\theta$ while all remaining terms change only by $O(\theta)$ on compacts (by continuity of the SSPE correspondence).

Consequently, for α_i large we can choose θ so small that $p^{(\theta)} \in \mathcal{U}$ and yet $\rho'_{l,i-1}[s_0]$ exceeds the threshold in Lemma 9 with $\varepsilon = 1/r$ (recall $\rho'_{l,i-1} = (2/(\kappa\alpha_{i-1}))(1 + \hat{\delta}V'_{l,i}) - 1$). Lemma 9 then yields $t_0 \in [\underline{t}, \bar{t}]$ such that $\rho^{-1}_{l,i-1}[\lambda_k t_0] - \rho^{-1}_{l,i-1}[t_0] < 1/r$. Hence $p^{(\theta)} \in \mathcal{O}_r \cap \mathcal{U}$.

Because $\mathcal{U}$ was arbitrary, $\mathcal{O}_r$ is dense. Finally, if $p \in \bigcap_{r\geq 1} \mathcal{O}_r$, then for every r there exist arbitrarily large m and interior layers at which (INC) fails with $\varepsilon = 1/r$; thus no uniform $\varepsilon > 0$ as in Condition 7 can exist. ■

Consequence. Proposition 8 should be read as a *conditional* statement: if the increment property (INC) were to hold uniformly on interior layers, then knife-edge perceived feasibility would force empty matching on compacts. Proposition 10 shows that (INC) is generically violated in PA (open and dense primitives), so this Paradox-of-Plenty mechanism is not robust.

Recall that in PA, actual feasibility implies perceived feasibility pointwise (actual $\subseteq$ perceived). Steepening of ρ can therefore not create agents who think they have no chance while actually having one; instead it segments the type space by shrinking the perceived overlap for many types, while generic failure of (INC) prevents the actual overlap from vanishing uniformly on compacts, leaving pockets of types with (small) actual—and hence also perceived—overlap.

So, in PA there are three types: 1) types with perceived and actual overlap; 2) those with perceived overlap but no actual overlap; 3) those with neither perceived nor actual overlap.

4.4 False hopers and type-space segmentation in partial awareness

The distinction between *perceived* versus *actual* mutual acceptance is useful for organizing who keeps searching (and why) under partial awareness away from the terminal layer.

Perceived versus actual overlap. Fix a side k (opponent $l \neq k$), a macro layer i, and a market type t_k with perceived type $s_k = \lambda_k t_k$. Write the *perceived* and *actual* widths of the mutual acceptance interval as

$$W^{\mathrm{perc}}_{k,i}(t_k) \;\equiv\; \frac{1}{\lambda_l}\,\rho^{-1}_{l,i}[\,\lambda_k t_k\,] \;-\; \rho_{k,i}[\,\lambda_k t_k\,], \qquad W^{\mathrm{act}}_{k,i}(t_k) \;\equiv\; \frac{1}{\lambda_l}\,\rho^{-1}_{l,i}[\,t_k\,] \;-\; \rho_{k,i}[\,\lambda_k t_k\,]. \quad (13)$$

A type t_k *perceives* a nonempty set of mutually acceptable partners at layer i iff $W^{\mathrm{perc}}_{k,i}(t_k) > 0$, and has a nonempty set of *actually* mutually acceptable partners iff $W^{\mathrm{act}}_{k,i}(t_k) > 0$.

Three sets of types. To avoid overloading symbols, we use calligraphic letters for the induced type sets:

$$\mathcal{P}_{k,i} \;\equiv\; \big\{t_k :\; W^{\mathrm{perc}}_{k,i}(t_k) > 0\big\}, \qquad \mathcal{A}_{k,i} \;\equiv\; \big\{t_k :\; W^{\mathrm{act}}_{k,i}(t_k) > 0\big\}, \quad (14)$$

$$\mathcal{F}_{k,i} \;\equiv\; \big\{t_k :\; W^{\mathrm{perc}}_{k,i}(t_k) > 0 \geq W^{\mathrm{act}}_{k,i}(t_k)\big\} = \mathcal{P}_{k,i} \setminus \mathcal{A}_{k,i}, \qquad \mathcal{N}_{k,i} \;\equiv\; \big\{t_k :\; W^{\mathrm{perc}}_{k,i}(t_k) \leq 0\big\}. \quad (15)$$

Thus $\mathcal{A}_{k,i}$ are the "real chance" types (your type (1)), $\mathcal{F}_{k,i}$ are the "false hopers" (type (2)), and $\mathcal{N}_{k,i}$ are the "no hope" types (type (3)).

Nesting: actual overlap implies perceived overlap. The perceived set always (weakly) contains the actual set. Indeed,

$$W^{\mathrm{perc}}_{k,i}(t_k) - W^{\mathrm{act}}_{k,i}(t_k) = \frac{1}{\lambda_l}\Big(\rho^{-1}_{l,i}[\,\lambda_k t_k\,] - \rho^{-1}_{l,i}[\,t_k\,]\Big) \;\geq\; 0, \quad (16)$$

since $\rho^{-1}_{l,i}$ is increasing and $\lambda_k > 1$. Hence $\mathcal{A}_{k,i} \subseteq \mathcal{P}_{k,i}$ and $\mathcal{F}_{k,i}$ is precisely the "boundary layer" of types who perceive overlap but in fact have none.

Where these sets live away from terminal, for large m. On interior macro layers (away from the terminal layer) and on compacts bounded away from 0, Lemma 2.3 implies a quantitative knife-edge: whenever $t_k \in \mathcal{P}_{k,i}$, the perceived overlap is small,

$$W^{\mathrm{perc}}_{k,i}(t_k) = O(m^{-1/2}). \quad (17)$$

Consequently, both $\mathcal{P}_{k,i}$ and its subset $\mathcal{A}_{k,i}$ (if nonempty) correspond to types whose perceived (and actual) match probabilities are small on compacts.

The distinction between $\mathcal{A}_{k,i}$ and $\mathcal{F}_{k,i}$ is driven by the *misperception correction*

$$\Delta_{l,i}(t_k) \;\equiv\; \rho^{-1}_{l,i}[\,\lambda_k t_k\,] - \rho^{-1}_{l,i}[\,t_k\,], \qquad W^{\mathrm{act}}_{k,i}(t_k) = W^{\mathrm{perc}}_{k,i}(t_k) - \frac{1}{\lambda_l}\Delta_{l,i}(t_k). \quad (18)$$

Types in $\mathcal{F}_{k,i}$ are those for which the correction term $\Delta_{l,i}(t_k)$ is large relative to the already small knife-edge overlap (17), so that it wipes out the actual overlap. By contrast, Proposition 10 shows that on any fixed compact away from 0 and on interior layers, failure of (INC) is generic: one can find (robustly in primitives) points where $\Delta_{l,i}(t_k)$ is arbitrarily small. Near such points, perceived and actual overlap nearly coincide, so the surviving "pockets" of actual overlap live inside $\mathcal{P}_{k,i}$ but are concentrated where the increment $\Delta_{l,i}(\cdot)$ is small.

Optional shape refinement (interval picture). If, on the relevant band, both $\rho_{k,i}$ and $\rho_{l,i}$ are convex, then $\rho_{l,i}^{-1}$ is concave and the widths $W_{k,i}^{\text{perc}}(\cdot)$ and $W_{k,i}^{\text{act}}(\cdot)$ inherit a single-crossing/concavity structure. In that case $\mathcal{P}_{k,i}$ and $\mathcal{A}_{k,i}$ are intervals (typically tail segments), and the false hopers form the interval difference $\mathcal{F}_{k,i} = \mathcal{P}_{k,i} \setminus \mathcal{A}_{k,i}$. This refinement is not needed for the genericity logic, but it provides a useful mental picture: away from terminal, most types fall in $\mathcal{N}_{k,i}$, a thin band $\mathcal{P}_{k,i}$ perceives a chance, and within that band a thinner sub-band $\mathcal{A}_{k,i}$ actually can match.

4.5 PA Example

This example illustrates how, under *Partial Awareness* (PA), *false hope* becomes the dominant outcome away from terminal as the subperiod resolution m increases.

Calibration. We impose symmetry $\lambda_A = \lambda_B = \lambda$ and use a lognormal market distribution $T \sim \text{LN}(\mu, \sigma^2)$ with mean normalized to one (so $\mu = -\frac{1}{2}\sigma^2$). The baseline parameters are

$$\lambda = 1.1, \qquad \kappa = 1.2, \qquad \delta = 0.9, \qquad \sigma = 0.6, \qquad n = 1,$$

so the effective horizon is $N = m$. We evaluate *away from the terminal layer* at the interior subperiod $i = i(m) \equiv \lfloor m/2 \rfloor$. To focus on a compact subset of types that is bounded away from zero and contains most of the lognormal mass, we report results on

$$K \equiv \left[\exp(\mu - 3\sigma),\, \exp(\mu + 3\sigma)\right] \subset (0, \infty).$$

Perceived vs. actual overlap. Given the equilibrium cutoff $\rho_i[\cdot]$ on the interior layer, define the perceived and actual *mutual-acceptance widths*

$$W_i^{\text{perc}}(t) \equiv \frac{1}{\lambda}\rho_i^{-1}(\lambda t) - \rho_i(\lambda t), \qquad W_i^{\text{act}}(t) \equiv \frac{1}{\lambda}\rho_i^{-1}(t) - \rho_i(\lambda t).$$

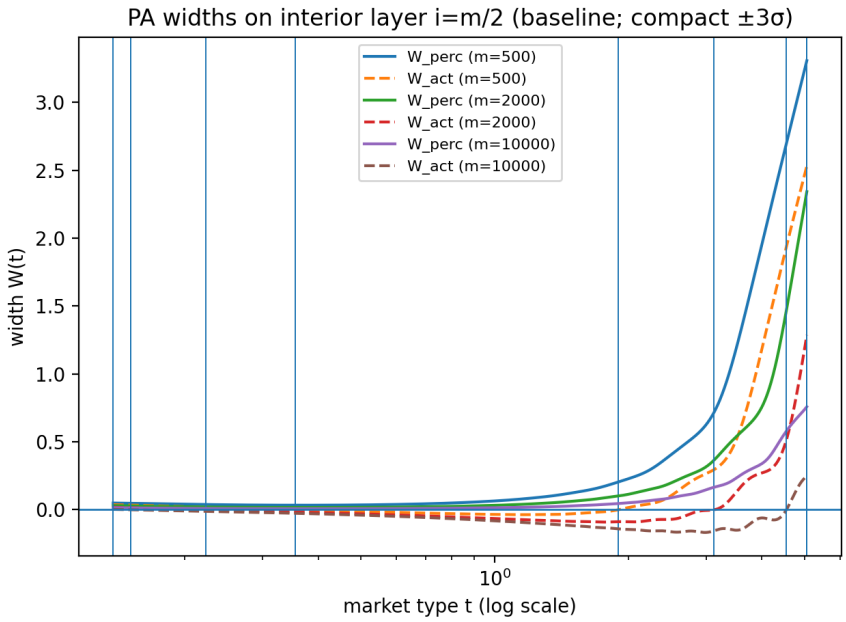


Figure 1: PA example (baseline): perceived and actual mutual-acceptance widths on the interior layer $i = \lfloor m/2 \rfloor$ for $m \in \{500, 2000, 10000\}$, on $K = [\exp(\mu - 3\sigma), \exp(\mu + 3\sigma)]$. Solid: W_i^{perc}; dashed: W_i^{act}. Vertical markers indicate the boundaries of $\mathcal{A}_i$.

These widths are the lengths (in opponent-type space) of the *perceived* and *actual* mutual acceptance intervals for market type t. Accordingly, on K we classify types into the four sets

$$\mathcal{P}_i = \{t \in K : W_i^{\text{perc}}(t) > 0\}, \quad \mathcal{A}_i = \{t \in K : W_i^{\text{act}}(t) > 0\}, \quad \mathcal{F}_i = \mathcal{P}_i \backslash \mathcal{A}_i, \quad \mathcal{N}_i = \{t \in K : W_i^{\text{perc}}(t) \leq 0\}.$$

Thus $\mathcal{F}_i$ is the set of *false hopers*: types who believe mutual acceptance is feasible but for whom it is not.

Mechanism. Figure **??** plots $W_i^{\text{perc}}(\cdot)$ and $W_i^{\text{act}}(\cdot)$ on K for $m \in \{500, 2000, 10000\}$ at $i = \lfloor m/2 \rfloor$ (dashed: actual, solid: perceived). Two features stand out. First, perceived overlap remains strictly positive throughout K (so $\mathcal{P}_i = K$ and $\mathcal{N}_i = \varnothing$ on this compact). Second, actual overlap collapses: $W_i^{\text{act}}(\cdot)$ becomes negative on most of K and remains positive only in shrinking extreme bands. Hence $\mathcal{A}_i$ becomes small while $\mathcal{F}_i$ expands.

False hope dominates as m grows. Figure **??** reports, on K, the probability mass shares of $\mathcal{A}_i$ (actual overlap), $\mathcal{F}_i$ (false hopers), and $\mathcal{N}_i$ (perceived empty) as functions of m at $i = \lfloor m/2 \rfloor$. In this baseline, $\mathcal{N}_i$ remains absent on K, but $\mathcal{F}_i$ quickly dominates: by

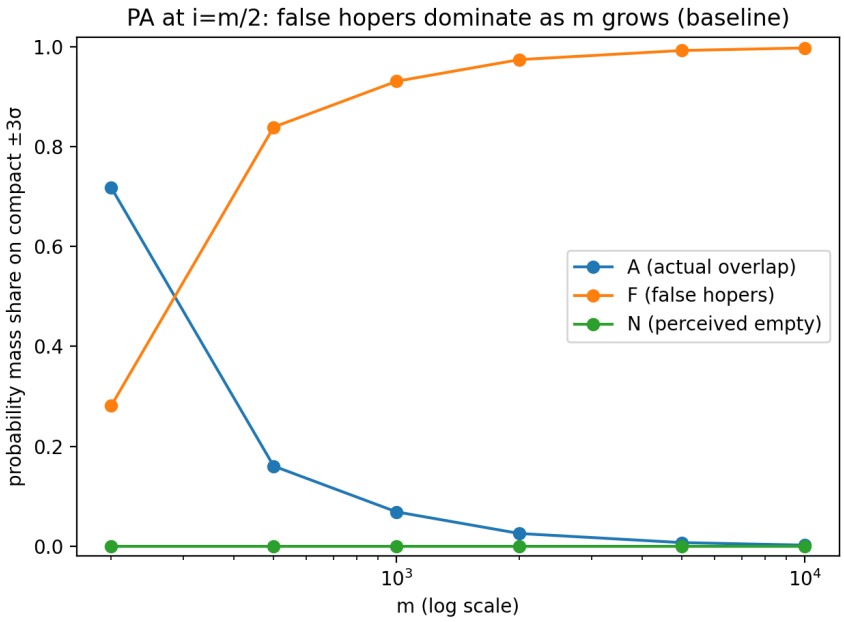


Figure 2:

$m = 2000$ more than 97% of types in K are false hopers, and by $m = 10000$ the false-hoper share exceeds 99.8%.

Figure **??** repeats the mass-share plot on a wider domain $\big[\exp(\mu - 6\sigma), \exp(\mu + 6\sigma)\big]$. In this calibration, perceived overlap remains positive even on that wider band, so $\mathcal{N}_i$ remains negligible; the main message is unchanged: increasing the draw resolution m causes false hope to dominate away from terminal.

A Appendix

A.1 Monotonicity.

Lemma 11 (Monotonicity in s) *In all three cases, UA, FA, and PA, the threshold function $\rho_{k,i}[\cdot]$ is strictly increasing.*

Proof. Fix i and let $s' > s$, with $\varepsilon \equiv s' - s > 0$. Take any measurable acceptance policy for a type with perceived state s in subperiod i. We construct a feasible policy for perceived state s' as follows: after every history, the s'–type takes the same accept/reject decision as the s–type, *except* that it rejects any draw t_l for which the opponent would reject a k-partner with

m	$\mathbb{P}(T \in \mathcal{A}_i \mid T \in K)$	$\mathbb{P}(T \in \mathcal{F}_i \mid T \in K)$	$\mathbb{P}(T \in \mathcal{N}_i \mid T \in K)$
200	0.718	0.281	0.000
500	0.160	0.839	0.000
1000	0.069	0.931	0.000
2000	0.025	0.975	0.000
5000	0.007	0.993	0.000
10000	0.002	0.998	0.000

Table 1: PA example (baseline): mass shares on K at the interior layer $i = \lfloor m/2 \rfloor$.

perceived state s after that history. (This is well defined because the opponent's acceptance decision is of cutoff form in the partner's type.) Because the opponent's cutoff acceptance is weakly increasing in the partner's market type, and because $s' > s$ corresponds to a (weakly) higher market type for player k (trivially in the full-unawareness case and via $s_k = \lambda_k t_k$ with $\lambda_k > 0$ in the FA cases), any opponent who would accept a k-partner with perceived state s must also accept one with perceived state s'. This construction therefore ensures that, under the same realized sequence of meetings, a match forms for type s' if and only if it forms for type s, and in that event the match time and partner type coincide.

Along any realized history, either both types remain single or both types match with the same partner t_l. In the single case, the flow payoff difference is $s' - s = \varepsilon \geq \frac{\kappa}{2}\varepsilon$ (since $\kappa < 2$); in the matched case, the flow payoff difference is exactly $\frac{\kappa}{2}(s' + t_l) - \frac{\kappa}{2}(s + t_l) = \frac{\kappa}{2}\varepsilon$. Thus the discounted continuation payoff from subperiod i for type s' under the constructed policy exceeds that for type s by at least

$$V_{k,i}[s'] - V_{k,i}[s] \geq \frac{\kappa}{2}\alpha_i \varepsilon. \tag{19}$$

where $\alpha_i \equiv \sum_{h=0}^{H_i - 1} \hat{\delta}^h = \frac{1 - \hat{\delta}^{H_i}}{1 - \hat{\delta}}$.

Finally, apply the threshold formula for $\rho_{k,i}[\cdot]$ and subtract:

$$\rho_{k,i}[s'] - \rho_{k,i}[s] = \frac{1}{\alpha_i}\frac{2}{\kappa}\left(\varepsilon + \hat{\delta}\left(V_{k,i+1}[s'] - V_{k,i+1}[s]\right)\right) - \varepsilon.$$

Using (19) with i replaced by $i + 1$ (so $H_{i+1} = H_i - 1$) yields $V_{k,i+1}[s'] - V_{k,i+1}[s] \geq \frac{\kappa}{2}\alpha_{i+1}\varepsilon$, and substituting this bound implies

$$\rho_{k,i}[s'] - \rho_{k,i}[s] \geq \frac{1}{\alpha_i}\left(\frac{2}{\kappa} - 1\right)\varepsilon \;>\; 0.$$

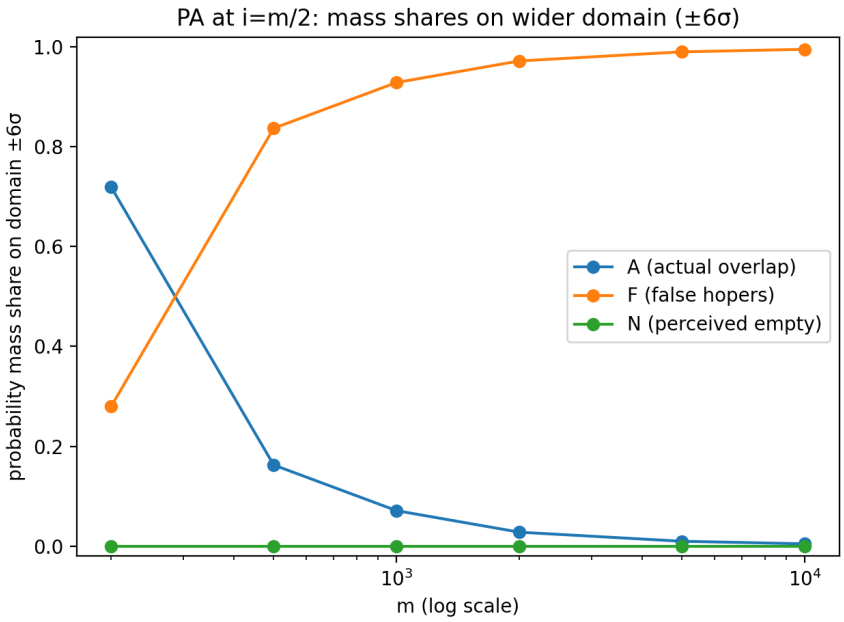


Figure 3: PA example (baseline): mass shares on the wider domain $\big[\exp(\mu-6\sigma), \exp(\mu+6\sigma)\big]$ at the interior layer $i = \lfloor m/2 \rfloor$.

Hence $\rho_{k,i}[\cdot]$ is strictly increasing. ■

Lemma 12 (Monotonicity i) *In all three cases, UA, FA, and PA, agents become less picky as the horizon shortens. Formally, for every s, $\rho_{k,i+1}[s] \leq \rho_{k,i}[s]$ for all $i < N$.*

Moreover, the inequality is strict whenever, starting from $(i+1, s)$ in the agent's subjective problem, there is positive probability of encountering before the horizon ends a draw for which mutual consent is (subjectively) feasible and $\frac{\kappa}{2}(s + \tilde{t}) > s$ (i.e., a strictly profitable match relative to remaining single).

Proof. Fix k and $s > 0$. Recall that $\alpha_i \equiv \sum_{h=0}^{H_i-1} \hat{\delta}^h$ denotes the annuity factor from subperiod i (so $\alpha_i = 1 + \hat{\delta}\alpha_{i+1}$ and $\alpha_i > \alpha_{i+1}$). In all three cases, conditional on a draw of partner market value t, if the opponent would not consent when player k consents, then k's consent is payoff–equivalent to rejection (no match forms and the agent continues). Hence we may restrict attention to draws for which mutual consent is feasible. If a match forms with partner market value t in subperiod i, the present value of the match payoff is $\alpha_i \frac{\kappa}{2}(s+t)$; if the agent rejects (or if no match forms), his continuation value is $s + \hat{\delta} V_{k,i+1}[s]$.

Define the normalized option value (relative to staying single)

$$w_{k,i}[s] \equiv \frac{V_{k,i}[s] - \alpha_i s}{\alpha_i},$$

and the normalized continuation weight

$$\beta_i \equiv \frac{\hat{\delta}\alpha_{i+1}}{\alpha_i} = 1 - \frac{1}{\alpha_i}.$$

Finally, define the (per-period) normalized gain from matching rather than staying single,

$$d\,[s,t] \equiv \frac{\kappa}{2}(s+t) - s,$$

which is strictly increasing in t.

Weak Increasingness

Step 1 (normalize the acceptance criterion). Using $\alpha_i = 1 + \hat{\delta}\alpha_{i+1}$, we can rewrite the rejection payoff as

$$s + \hat{\delta}V_{k,i+1}[s] = s + \hat{\delta}\alpha_{i+1}\left(s + w_{k,i+1}[s]\right) = \alpha_i s + \alpha_i\beta_i w_{k,i+1}[s].$$

Subtracting the single baseline $\alpha_i s$ from both match and rejection payoffs and dividing by α_i yields the equivalent acceptance rule

$$d\,[s,t] \geq \beta_i\, w_{k,i+1}[s].$$

Whenever matching is feasible at (i,s), the cutoff $\rho_{k,i}[s]$ is therefore characterized by

$$d\left[s, \rho_{k,i}[s]\right] = \beta_i\, w_{k,i+1}[s].$$

(If matching is infeasible at (i,s), any cutoff is payoff–equivalent; the statement of the lemma is then immediate.)

Step 2 (the option value of waiting is weakly larger earlier). The normalized value satisfies the recursion

$$w_{k,i}[s] = \mathbb{E}\big[\max\left\{d\left[s,\tilde{t}\right],\, \beta_i\, w_{k,i+1}[s]\right\}\big], \qquad w_{k,N}[s] = \mathbb{E}\big[\max\left\{d\left[s,\tilde{t}\right],\, 0\right\}\big],$$

where $\tilde{t} \sim F$. Here, on draws $\tilde{t}$ for which a match would not form even if player k consents, the continuation term applies, so such draws do not affect the acceptance cutoff. Since

$\alpha_i > \alpha_{i+1}$, we have $\beta_{i+1} \leq \beta_i$ for all $i < N$. We claim that $w_{k,i}[s] \geq w_{k,i+1}[s]$ for all $i < N$. For $i = N-1$, $w_{k,N}[s] \geq 0$ implies

$$w_{k,N-1}[s] = \mathbb{E}\big[\max\{d\left[s,\tilde{t}\right], \beta_{N-1} w_{k,N}[s]\}\big] \geq \mathbb{E}\big[\max\{d\left[s,\tilde{t}\right], 0\}\big] = w_{k,N}[s].$$

For the induction step, assume $w_{k,i+1}[s] \geq w_{k,i+2}[s]$. Since $\max\{a, \cdot\}$ is nondecreasing in its second argument and $\beta_{i+1} \leq \beta_i$,

$$w_{k,i+1}[s] = \mathbb{E}\big[\max\{d\left[s,\tilde{t}\right], \beta_{i+1} w_{k,i+2}[s]\}\big] \leq \mathbb{E}\big[\max\{d\left[s,\tilde{t}\right], \beta_{i+1} w_{k,i+1}[s]\}\big] \leq \mathbb{E}\big[\max\{d\left[s,\tilde{t}\right], \beta_i w_{k,i+1}[s]\}$$

Step 3 (cutoffs fall over time). Combining $w_{k,i+1}[s] \geq w_{k,i+2}[s]$ with $\beta_{i+1} \leq \beta_i$ gives

$$\beta_{i+1} w_{k,i+2}[s] \leq \beta_i w_{k,i+1}[s].$$

Because $d(s,t)$ is strictly increasing in t, a weakly smaller right-hand side expands the acceptance set. Hence the reservation type solving $d\left[s, \rho_{k,i}[s]\right] = \beta_i w_{k,i+1}[s]$ is weakly decreasing in i: $\rho_{k,i+1}[s] \leq \rho_{k,i}[s]$ for all $i < N$.

Strict increasingness

Fix $i < N$ and $s > 0$. Whenever matching is (subjectively) feasible at both (i,s) and $(i+1,s)$, Step 1 above gives the cutoff characterizations

$$d\big[s, \rho_{k,i}[s]\big] = \beta_i\, w_{k,i+1}[s], \qquad d\big[s, \rho_{k,i+1}[s]\big] = \beta_{i+1}\, w_{k,i+2}[s],$$

with $d(s,\cdot)$ strictly increasing. By Step 2, $w_{k,i+1}[s] \geq w_{k,i+2}[s] \geq 0$, and since $\alpha_i > \alpha_{i+1}$ we have $\beta_i > \beta_{i+1}$. Hence

$$\beta_i\, w_{k,i+1}[s] \;\geq\; \beta_{i+1}\, w_{k,i+2}[s],$$

which implies $\rho_{k,i+1}[s] \leq \rho_{k,i}[s]$. Moreover, if $w_{k,i+1}[s] > 0$, then $\beta_i\, w_{k,i+1}[s] > \beta_{i+1}\, w_{k,i+2}[s]$, hence $\rho_{k,i+1}[s] < \rho_{k,i}[s]$.

Finally, by the recursion for $w_{k,i+1}[s]$ in Step 2,

$$w_{k,i+1}[s] = \mathbb{E}\big[\max\left\{d\left[s,\tilde{t}\right], \beta_{i+1} w_{k,i+2}[s]\right\}\big],$$

so $w_{k,i+1}[s] > 0$ holds exactly when, under the agent's subjective beliefs at state s (starting from $i+1$), there is positive probability of encountering before the horizon ends a draw for which mutual consent is (subjectively) feasible and $d[s,\tilde{t}] > 0$, i.e. a match that yields a strictly higher flow payoff than remaining single.

This establishes the lemma. ■