Building and Maintaining Productive Relationships*

Nemanja Antić[†]

Ameet Morjaria[‡]

Miguel Ángel Talamas Marcos[§]

November 2025

Productive relationships require a mutual understanding of what each party requires. We model relationship formation and termination when agents face uncertainty about which actions (if any) are productive, what we call a "clarity problem." In each period of experimentation, the agent finds a productive action with some probability; this sequence of probabilities describes clarity. Our model predicts a reversal in termination rates: agents with higher clarity are less likely to end relationships in the early search phase, but more likely to end them later, since they find productive relationships easier to replace. Using transaction-level data from Ethiopia's rose industry, we document that clarity problems are economically significant. Moreover, exporters that are more likely to form productive relationships, those with higher clarity, are also more likely to terminate them when outside options improve. This explains why domestically owned exporters, despite lower discount factors, are less responsive to positive shocks to the outside option than foreign-owned exporters. Clarity emerges as a key determinant of relationship formation and maintenance, determining whether exporters can access the higher profitability and economic growth associated with direct sales of differentiated products.

JEL classifications: D21, D22, D83, D86, L14

*We are especially grateful to Charles Angelucci, Sandeep Baliga, Dan Barron, Alessandro Bonatti, Jeff Ely, Bob Gibbons, Thomas Jungbauer, Jin Li, Jim Malcomson, Niko Matouschek, Meg Meyer, Joris Mueller, Harry Pei, Michael Powell, Doron Ravid, Luis Rayo, Bruno Strulovici, Alireza Tahbaz-Salehi, Marta Troya-Martinez, and Diego Vera for numerous helpful discussions. We thank seminar and conference participants at ASSA, Cornell, Transatlantic Theory Workshop 2024, George Mason, IADB, MSU, NEUDC, Northwestern, NBER/AfDB Transforming Rural Africa, Tecnológico de Monterrey, EGAP CDMX, Colmex, Makris Symposium in Economic Theory, SIOE, EPED Montreal, and the 7^{th} Relational Contracting Conference for many comments and suggestions. For excellent research assistance we thank Maria Emilia Bullano, Alejandra Goytia and Alison Zhao. We thank the Ethiopian Investment Agency (EIA), Ethiopian Horticultural Development Agency (EHDA), Ethiopian Horticulture Producer Association (EHPA), Ministry of Trade & Industry (MoTI), the Development Bank of Ethiopia (DBE) and many exporters, especially Ato Joseph Iskander and Ato Tsegaye Abebe, for data and sharing industry knowledge. We are grateful to the National Graduate Institute for Policy Studies, Foundation for Advanced Studies on International Development, and the Ethiopian Development Research Institute and in particular to Tomoya Matsumoto and Yuki Mano for sharing data. Morjaria thanks the EHDA, EHPA, MoTI and especially EDRI and Ato Getachew Yoseph for hosting countless visits over the last ten years. Financial support from the CIFAR Azrieli Global Scholars Program, GPRL at Northwestern and Ford Motor Center for Global Citizenship at Kellogg School of Management is gratefully acknowledged. The findings, interpretations, and conclusions expressed herein are entirely those of the authors and do not necessarily reflect the views of any organizations and individuals that shared information.

[†]Northwestern University, nemanja.antic@kellogg.northwestern.edu

[‡]MIT Sloan, CEPR & NBER morjaria@mit.edu

[§]Inter-American Development Bank, miguelta@iadb.org

1 Introduction

When contracts are incomplete or externally unenforceable, the value of future interactions can give rise to self-enforcing relationships (e.g., Macleod and Malcomson, 1989; Baker et al., 1994, 2002; Levin, 2003). But successful relationships require more than proper incentives. Parties must clearly communicate their needs and understand what the other expects, which can be difficult. Such clarity problems are central, for example, when a user interacts with a large language model: it is miscommunication, not incentive misalignment, that typically prevents productive outcomes. These clarity problems are common in standard economic settings, such as employer-employee or supplier-buyer interactions, since describing, discerning, and delivering on expectations in a relationship is not straightforward (Gibbons and Henderson, 2012).

Consider a grader working for a professor. The professor may know the outcomes she wants (appropriate partial credit, the right tone and detail of comments, how closely scores should match a target distribution), yet communicating and comprehending these expectations is difficult. Rubrics help, but are incomplete due to practical limits or unforeseeable boundary cases, so clarity problems persist. As a result, the grader must use his judgment and may fail to meet the professor's expectations. Failures can be informative: discovering that partial credit awarded was beyond what the professor finds acceptable narrows the set of actions to try next time. However, if such misunderstandings are not resolved, the relationship will break down. Similar problems arise in employer-employee relationships where workers must acquire firm-specific human capital (Lazear, 2009), or even in romantic relationships.

Clarity problems also arise in supplier-buyer relationships. Take rose exports, our empirical context. Because roses are fragile and highly perishable, precise logistics are critical for these relationships to be productive. Exporters must understand and adapt shipments to each buyer's circumstances, including product mix, bloom ripeness, packaging, flower food and transit microclimate. While buyers may know what condition they want the roses to arrive in, this does not easily translate into specific actions an exporter must take. These clarity problems are consequential: roughly half of all relationships fail within the first few shipments. Such failures are economically significant, as exporting differentiated products directly to buyers (rather than commodities through auction markets) is linked to higher seller profitability (Cajal-Grossi et al., 2023) and economic growth (e.g., Rauch, 1999; Atkin et al., 2024).

In this paper, we study how clarity problems influence relationship formation and termination. We develop a principal–agent model in which parties must resolve a clarity problem before their relationship can become productive. Our main theoretical result is a reversal in termination probabilities: agents with higher clarity (those better at discovering productive actions) are less likely to terminate a relationship initially than those with lower clarity, but more likely to do so once the relationship is productive. This reversal does not arise in canonical models of learning match value like Jovanovic (1979), where we show that better match quality reduces termination probabilities across all time periods.[1] We empirically investigate our theoretical

[1]This is because in Jovanovic (1979) match quality affects the agent's wage in each period. In our model, clarity matters only for finding productive actions, but not for the surplus they generate once found. This makes

predictions using transaction-level data from Ethiopian flower exports, where we document that clarity problems are substantial. We find evidence consistent with this reversal: exporters with higher estimated clarity terminate relationships less early on, but respond more strongly to outside option shocks once relationships become productive.

In our model, a long-lived agent must resolve a clarity problem each time he starts a new relationship with a principal. Specifically, there is uncertainty about which of the agent's actions, if any, will be productive and generate a positive surplus for the principal. In the initial search phase of a relationship, the agent experiments with different actions. The principal observes the chosen action, and decides whether to pay a fixed bonus, which signals to the agent that the action was productive. Instead of trying a new action, the agent can terminate the relationship by taking a stochastic spot payoff (e.g., the exporter can sell at the auction in our empirical application). At the end of each period, the agent may pay a cost to be matched to a new principal. If the agent finds a productive action, the relationship transitions to its productive phase, where the agent can keep taking that action as long as he prefers it to his outside option.

During the search phase of a relationship, we assume that the probability λ_t of successfully finding a productive action is single-peaked and converges to zero over time, t.[2] This allows the probability of success to rise initially, as early failures can be informative about which actions are likely to be productive. Eventually, however, persistent failures make the agent pessimistic about the possibility of establishing a productive relationship with the current principal. We interpret the sequence of λ_t as *clarity*, which reflects both players' characteristics, as well as the complexity of the underlying environment.

When there are no clarity problems, $\lambda_0 = 1$, the agent immediately knows a productive action for the principal. When $\lambda_0 < 1$ there is a clarity problem in the relationship, which is resolved only if a productive action is found. If the agent learns a productive action for one principal, it does not change his beliefs about what actions are productive for another principal, since principals have different preferences. Therefore, each new relationship faces the clarity problem and starts from the search phase. We characterize the agent's optimal strategy under the assumption that the principal pays the bonus whenever the agent takes a productive action.[3]

The agent remains in a relationship as long as the expected benefit of doing so exceeds his outside option. The outside option has two components: a stochastic spot payoff and the continuation value of starting a new relationship. The agent ends a relationship if the realized spot payoff exceeds a threshold. This threshold is higher once a productive action has been found, so productive relationships are less likely to end than those in the search phase.

We derive comparative statics on how clarity affects relationship termination. Our model shows that higher clarity, i.e., increasing any λ_t, raises the threshold for ending a relationship up to period t, but lowers it afterward. This implies a reversal in termination rates: higher clarity reduces early relationship termination, but increases it in later periods, especially once only productive relationships remain. The reversal arises because, in the productive phase of

productive relationships more replaceable than in Jovanovic (1979).

[2]Although we provide a microfoundation for this assumption in the context of the Pandora's box problem with correlation (Weitzman, 1979), it is helpful to think of the sequence (λ_t) as exogenous.

[3]By focusing on the agent's dynamic programming problem, we isolate the core trade-offs in the main text. Section 3.4.2 describes how principal incentives can be added to the model and shows that principals may strategically avoid high clarity agents, deeming them "overqualified".

a relationship, higher clarity increases the agent's value of starting a new relationship, making the current relationship more replaceable. In the search phase, however, the direct effect of higher clarity on the continuation value of the current relationship exceeds its indirect effect on the value of a new one, making the agent less likely to take the outside option.

Testing these theoretical predictions requires us to address measurement and causal identification challenges. It requires a context where limited external enforcement makes relationships critical, along with transaction-level data with buyer and seller identifiers to track relationships from beginning to end. The central identification challenge is that relationship termination is endogenous: exporter, buyer, and relationship characteristics influence survival. Therefore, any effort to identify a causal effect requires exogenous variation in termination probability, ideally through observable outside options that affect this likelihood.

To overcome measurement challenges, we utilize a decade of transaction-level data with exporter and buyer identifiers from the Ethiopian rose industry, supplemented by two exporter censuses. Exporters trade through two channels: the spot market auction and direct relationships with international buyers. The auction is a natural outside option since exporters can redirect shipments from relationships to the auction at any time. On average, exporters sell roughly half of their production through each channel, however, domestically owned exporters transact considerably less through relationships than their foreign-owned counterparts. The two groups are otherwise similar in rose quality, production processes, and scale.

To address the identification challenge, we leverage exogenous variation in the auction price. On average, prices are higher in direct relationships because parties avoid intermediary costs, and shipments are tailored to each buyer's requirements. However, since prices in direct relationships are set at the beginning of the season, while auction prices fluctuate, auction prices sometimes exceed relationship prices. Since Ethiopian exporters are a small fraction of the global market, these price fluctuations are likely exogenous to them.

Consistent with our model's predictions in the presence of clarity problems, the probability of relationship termination is high during the first four periods (shipments) and stabilizes thereafter. Domestically owned firms are particularly likely to end relationships early, suggesting they face more severe clarity problems, i.e., they have lower clarity. In many standard models, including ours, a lower discount factor increases the likelihood that the agent ends a relationship in favor of a short-term gain.[4] However, in our model, lower clarity can counteract this effect if the discount factor and clarity are positively correlated, as is the case in our empirical context. We find that, despite a lower discount factor (higher cost of capital), domestically owned exporters, who also have lower clarity, are less likely to end a relationship in response to positive shocks to the auction price.

The lower clarity of domestically owned firms could be attributed to themselves or to their environment. For example, if they were facing more demanding buyers, it would be harder for them to find productive actions. A better understanding of the source of clarity problems is important for designing policies to address them. In Section 3.4.1, we provide a microfoundation for the model where clarity is decomposed into an exporter and a buyer component. In our data, buyers interact with multiple exporters and vice versa, so we can estimate the buyer and

[4] See Quah and Strulovici (2013) for general conditions under which this comparative static obtains.

exporter components of clarity using the AKM framework (Abowd et al., 1999). We find that both components are crucial for reaching a productive relationship. Moreover, we find that domestically owned firms have a lower exporter component, but do not face worse buyers.

Finally, we test the key prediction in the model: the reversal in termination rates. Exporters with an estimated higher clarity component, which are less likely to terminate relationships early, are more likely to end relationships later (in their productive phase) in response to positive shocks to the auction price.

Related Literature. We contribute to a growing literature on how agents form and sustain cooperative relationships. We are not the first to highlight the importance of "clarity". Gibbons and Henderson (2012) offer a broader interpretation of clarity, encompassing the possibility that players may not understand the equilibrium they are expected to play. Our definition of clarity is more narrow: we focus on incomplete information about which actions are productive. Chassang (2010) also studies a setting where only some actions are productive, but the agent knows which ones are, while the principal does not. In his framework, learning is intertwined with incentive problems because the informed agent is tempted to take unproductive actions, while the uninformed principal must sometimes use inefficient on-path punishments, and it may be optimal to stop learning even when more efficient routines exist. In our model, by contrast, the agent does not know which actions are productive (and the principal cannot communicate this even if she knew it). We also add a stochastic outside option with rematching, enabling us to study how clarity and outside opportunities jointly shape relationship formation and termination. Closer to how we model learning within relationships, Bonatti and Hörner (2011) study team-production with hidden effort where free-riding leads to procrastination.

We contribute to a well-established literature that examines informal agreements, such as relational contracts (Macleod and Malcomson, 1989; Baker et al., 1994, 2002; Levin, 2003), and in particular models with incomplete information about the value of the relationship (Halac, 2012; Kostadinov and Kuvalekar, 2022).[5] In Fuchs (2007), Deb et al. (2016) and Fong and Li (2016) subjective or privately observed performance signals interact with relational incentives, but uncertainty is about types or performance, rather than about which actions are productive. Models of relationship formation in this literature emphasize the role of reputation instead of clarity problems (e.g., Ghosh and Ray, 1996; Kranton, 1996; Watson, 1999).[6] These papers find that parties increase the scale of a relationship over time, which is not something we observe in our empirical context. McAdams (2011) studies a general model of stochastic partnerships, but does not consider the kind of comparative statics that are central to our theoretical and empirical contributions. Board and Meyer-ter Vehn (2015) study relational contracting with rematching in a competitive labor market and focus on how identical firms may offer different relational contracts to workers.

Our model closely relates to the labor search literature on learning about match quality (Jovanovic, 1979, 1984; Moscarini, 2005).[7] However, Jovanovic (1979) does not provide comparative

[5] For reviews, see MacLeod (2007), Malcomson (2013), and Fahn et al. (2023).

[6] Halac (2014) and Angelucci and Orzach (2025) study relationship dynamics when players can choose the project they work on.

[7] See Mortensen and Pissarides (1999) for a survey of this part of the literature, and Wright et al. (2021) for a survey which focuses on directed search.

statics on how the distribution of match quality (which relates to clarity) affects termination rates. We derive this comparative static in Online Appendix B and show that the reversal in termination rates does not arise in his framework. This difference stems from how we model productivity. In our model, the surplus from a productive relationship is known ex-ante, but players must learn whether such a relationship is possible. In contrast, in Jovanovic (1979), a productive relationship is always possible, but its surplus is uncertain.[8] As a result, agents who draw from a better distribution of match quality in Jovanovic's model are less likely to terminate relationships in both early and late periods, whereas in our model, high-clarity agents exhibit lower early termination but higher late termination.[9]

Our paper relates to the broader literature on search and matching (Chade et al., 2017), especially learning while searching (Adam, 2001). Our search microfoundation in Appendix A considers the Pandora's box problem (Weitzman, 1979) with correlation. This problem is known to be intractable in general,[10] but we depart from the literature in two ways. First, the structure of our problem (prizes are binary) implies a trivial stopping rule. Second, we are interested in the conditional probability of success (as this feeds into our notion of clarity), which has not been a focus of abstract search models. In this respect, work on consumer search on platforms is more closely related. Nocke and Rey (2024) develop a model of sequential consumer search in which products differ in their ex-ante probability of being an acceptable match and when they are randomly ordered by a platform, every failed inspection makes it more likely that the remaining products are more likely to be an acceptable match. This raises the perceived probability of eventually finding a good match as search continues and relates to the case where clarity is monotonically increasing (up to some finite n) in our model.[11]

Our paper also contributes to the growing empirical literature on relationships and contracting (McMillan and Woodruff, 1999; Banerjee and Duflo, 2000). Examples of contemporary work in this literature include examining the role of seller market power and relational agreements (Brugués, 2024); adaptation (Barron et al., 2020; Gil et al., 2022); the interaction between relational contracts and spot market thickness (Harris and Nguyen, 2025a); and how buyer competition impacts relational contracts (Macchiavello and Morjaria, 2021); how long-term relationships through the threat of termination can deter opportunism (Harris and Nguyen, 2025b).[12] In contrast to these papers, our focus is on the ease of relationship formation and how it affects their strategic termination.

Similar to our empirical context, Macchiavello and Morjaria (2015) study relationships using data from Kenyan flower exporters, where they infer the value of relationships from exporters' willingness to forgo short-term gains. We instead study how relationships are built and show that when outside options are valuable, exporters who can more easily form new relationships

[8] In Jovanovic (1979), parties search for the most productive match possible. In our model, parties seek a match that is "good enough" (e.g., the exporter can deliver the roses the buyer needs, the TA can grade as the professor prefers), and there are potentially many such matches, making them more replaceable.

[9] In the same spirit, Kostadinov and Kuvalekar (2022) study a relational contracting model of learning about match quality. In their setting, the relational-contract surplus and likelihood of survival both increase with the belief that match quality is high, so no reversal emerges.

[10] Chawla et al. (2020) and Gergatsouli and Tzamos (2023) develop algorithms for (approximately) solving this problem, while Auster and Che (2025) study it with a regret-minimization objective.

[11] See Janssen et al. (2025) for search on platforms when match value is continuous as opposed to binary.

[12] For recent empirical reviews see Macchiavello (2022) and Macchiavello and Morjaria (2023).

are more likely to take them at the expense of losing existing relationships.

Lastly, there is growing evidence that exporting benefits firms, especially in low- and middle-income countries (e.g., Verhoogen, 2008; Atkin et al., 2017; Bastos et al., 2018; Hansman et al., 2020; Demir et al., 2024). Related work has examined how to induce firms to export and the effects of joining multinational supply chains (e.g., Volpe Martincus and Carballo, 2008; Alfaro-Ureña et al., 2022). Yet, beyond whether firms export at all, we emphasize the importance of how they export and whether they can form relationships with international buyers.

2 Model

An agent (e.g., an employee or exporter) can either take a stochastic spot payoff in each period or be matched to play a repeated game with a principal (e.g., a manager or buyer). All parties are risk-neutral, time is discrete, and the discount factor is $\delta < 1$. We describe the model using the employer-employee setting, but Section 4 details the exporter-buyer application that underlies our empirical analysis.

At the start of period t, the agent observes a spot payoff, $s_t \in [\ell, \infty) \subset \mathbb{R}_{++}$, drawn i.i.d from distribution F with mean $v = \int s \, \mathrm{d}F(s) < \infty$. Let A_t denote the agent's action set in period t. If the agent is unmatched, $A_t = \{s_t\}$, where s_t denotes the action of taking the spot option as well as the agent's payoff from it. The agent starts the game unmatched at $t = -1$.

There is a continuum of principals indexed i. If matched to principal i at time t, the agent's action set is $A_t = \{s_t\} \cup \mathcal{A}_i$, where actions in $\mathcal{A}_i$ correspond to working for principal i. If the agent chooses $a_t = s_t$, the relationship with i ends.[13] We say that the agent took the outside option, which gives him the spot payoff s_t today plus the (time-invariant) continuation value of rematching to another principal. We thus refer to s_t as the shock to the outside option. For simplicity, we assume that once a relationship ends, the agent cannot return to that principal.

Each action $a_k \in \mathcal{A}_i$ is either productive, generating a positive surplus for the principal, or non-productive, yielding no surplus. For example, a professor (principal) may only find certain ways of grading acceptable. Even if she provides a rubric, it may be incomplete, and so the grader (agent) will have to use his judgment. All productive actions generate the same surplus. Once a productive action is found, the relationship transitions from the search phase to the productive phase, where the agent can choose the same productive action in every period.

During the search phase, let λ_t denote the probability that the agent finds a productive action for the current principal on the t^{th} attempt, conditional on not finding one yet. We refer to the sequence $(\lambda_t)_{t=0}^{\infty}$ as *clarity*. Clarity is perfect only when $\lambda_0 = 1$ and the agent can immediately choose a productive action. Otherwise, the players face a *clarity problem* which is resolved through experimentation. Clarity problems arise either when the principal cannot communicate which actions are productive to the agent, or when she learns whether an action is productive only after the agent attempts it. We assume that clarity can initially improve as failures are informative, but that optimism eventually fades.

[13] We make this assumption because it is realistic in our empirical context, where the agent's outside option is high precisely when the value of a productive action is high for the principal, such as around Valentine's Day, when flower demand and auction prices peak.

Assumption 1. *The clarity sequence $(\lambda_t)_{t=0}^{\infty}$ is single-peaked, that is, there exists some $T \geq 0$ such that $(\lambda)_{t=0}^{T}$ is (weakly) increasing and $(\lambda_t)_{t=T}^{\infty}$ is (weakly) decreasing and converges to 0.*

The intuition behind Assumption 1 relies on discounting. If the probability that actions are productive is independent, the agent first tries the action that is most likely to be productive, in order to get the bonus as soon as possible and so $(\lambda_t)_{t=0}^{\infty}$ is (weakly) decreasing, i.e., single-peaked with a peak at $T = 0$. With correlation, he may opt for a different action, especially if that action is informative about what is likely to be productive, so that his chances of finding a productive action on the next attempt increase. Correlation can arise, for example, when a grader knows ex ante that exactly one of several ways of marking a question is acceptable to the professor. Assumption 1 allows for this possibility, under which failures can increase the probability of finding a productive action on the next attempt. Appendix A provides a microfoundation for this assumption in the context of the Pandora's box problem (Weitzman, 1979). We assume $\lambda_t \to 0$ to ensure that the agent will eventually terminate an unproductive relationship, rather than continuing to search forever.

Whenever the agent is matched to a new principal, he faces the same clarity problem, represented by $(\lambda_t)_{t=0}^{\infty}$.[14] So learning a productive action for principal i does not change the agent's beliefs about what action principal $j \neq i$ will find productive. The game effectively restarts whenever the agent is matched to a new principal and we refer to this period as $t = 0$.[15] Therefore, within a given relationship, t counts the number of unsuccessful attempts at finding a productive action. We can thus drop the index i and refer to a generic principal as *the* principal.

The principal privately observes her payoff and chooses $b_t \in \{0, b\}$. For now, we assume that the principal is committed to playing a truthful strategy: she pays the bonus if and only if a productive action was chosen.[16] At the end of each period, the agent chooses whether to pay a search cost $c \geq 0$ to be matched to a new principal in the next period. If the agent matches with a new principal, we reset time to $t = 0$.

In period t, the timing of actions is as follows:

1. The shock to the outside option $s_t \sim F$ is realized and observed by the agent.
2. The agent chooses an action $a_t \in A_t$.
3. If $a_t \neq s_t$, the principal observes her payoff and chooses $b_t \in \{0, b\}$.
4. The agent decides whether to pay cost $c \geq 0$ to match with a new principal next period.

The agent's stage game payoff in period t is $u : A_t \times \{0, b\} \to \mathbb{R}$, defined as follows: $u(s_t, \cdot) = s_t$, so if the agent chooses the outside option he gets the payoff s_t, and $u(a_t, b_t) = b_t$ for any $a_t \neq s_t$. We interpret $b > 0$ as the agent's per-period payoff when the principal pays the bonus. An agent who is indifferent between terminating a relationship and not breaks ties in favor of staying with the current principal.

[14] We discuss how our results extend to multiple principal types in Section 3.4.1.

[15] Thus, there is learning-by-doing (Bonatti and Hörner, 2011) within a relationship, but no learning across relationships. Every time the game restarts, new shocks are drawn.

[16] We show how the principal's incentives may be incorporated in Section 3.4.2.

We are interested in solving the agent's problem, that is, finding the optimal strategy for a long-lived agent. The agent begins unmatched at $t = -1$, takes the spot payoff and decides whether to pay matching cost $c \geq 0$. Since the agent could continue to take the spot payoff indefinitely, he will match to a principal only if

$$\delta W_0 - c \geq \frac{\delta \upsilon}{1-\delta}, \tag{RC}$$

where W_0 denotes the agent's continuation value if he starts period 0 matched to a new principal. If (RC) is not met, the agent always takes the spot payoff. In analyzing the agent's problem, we assume that this relationship constraint (RC) holds and then verify this condition ex post.

3 Theoretical Results

Let W_t denote the agent's continuation value at the start of period t, before shock s_t is drawn, while in the search phase of a relationship. After s_t is observed, define $W_t(s_t)$ as the agent's continuation value, so that $W_t = \mathbb{E}_F[W_t(s_t)]$. Let V be the agent's continuation value at the start of a period in the productive phase of a relationship, when a productive action is known.

In period t of the search phase, the agent either takes the outside option or works for the principal. The agent's value from taking the outside option is $s_t + \delta W_0 - c$, since when inequality (RC) holds the agent will pay cost $c \geq 0$ to rematch with a new principal in the following period. When working for the current principal in period t, the agent attempts an action that is productive with probability λ_t. If the action is productive, the agent receives a bonus b and the relationship transitions to the productive phase; otherwise, no bonus is paid and the relationship remains in the search phase. Thus the agent's value from working for the principal is

$$\lambda_t (b + \delta V) + (1 - \lambda_t) \max \{\delta W_{t+1}, \delta W_0 - c\}.$$

The first term is the benefit from resolving the clarity problem, while the second represents failure to find a productive action, after which the agent either continues with the same principal or restarts with a new one after paying cost $c \geq 0$. Conditional on shock s_t, the agent's continuation value in period t of the search phase is therefore

$$W_t(s_t) = \max \{s_t + \delta W_0 - c, \lambda_t (b + \delta V) + (1 - \lambda_t) \max \{\delta W_{t+1}, \delta W_0 - c\}\}. \tag{1}$$

By definition $V \geq W_t$ for all t, since in the productive phase the agent can always choose the productive action and receive the bonus for sure, as opposed to only being paid the bonus with probability λ_t in the search phase.

In the productive phase of a relationship, the agent either takes the productive action or the outside option. Taking the productive action results in an expected value of $b + \delta V$ for the agent. Taking the outside option instead yields $s_t + \delta W_0 - c$. Hence there exists a cutoff s^* such that the agent prefers to terminate a productive relationship if $s > s^*$, and we have that

$$V = \int_{\ell}^{s^*} (b + \delta V) \, \mathrm{d}F(s) + \int_{s^*}^{\infty} (s + \delta W_0 - c) \, \mathrm{d}F(s). \tag{2}$$

Given W_0 and V, s^* makes the agent indifferent between continuing and taking the outside option, so that

$$s^* = b + \delta V - \delta W_0 + c. \tag{3}$$

If $F(s^*) = 1$, no shock to the spot payoff will induce the agent to terminate a productive relationship and $V = \frac{b}{1-\delta}$.[17] Lemma 1 in Appendix B shows that for any W_0, there exists a unique V and s^*. Furthermore, V is increasing in W_0, while the total derivative of s^* with respect to W_0 is negative.

From equation (1), in each search period there is a cutoff shock, s_t^*, so that if $s_t > s_t^*$ the agent takes the outside option. This cutoff is

$$s_t^* = \lambda_t (b + \delta V) + (1 - \lambda_t) \max\{\delta W_{t+1}, \delta W_0 - c\} - \delta W_0 + c. \tag{4}$$

Because $b + \delta V \geq \max\{\delta W_{t+1}, \delta W_0 - c\}$, we have $s^* \geq s_t^*$ for all t, with strict inequality when $\lambda_t < 1$. Thus, holding fixed the shock distribution, productive relationships are (weakly) less likely to end than relationships in the search phase.

Theorem 1. *There is a unique solution to the agent's problem. The agent pays the search cost to match to a principal whenever inequality (RC) holds. A sufficient condition for inequality (RC) is*

$$\lambda_0 \geq \frac{(1-\delta)(\delta v + c)}{\delta b - \delta^2 v}.$$

Appendix B.1 shows the details of the existence and uniqueness proof, which does not rely on Assumption 1. A weaker sufficient condition for existence is given in the proof of the theorem, namely that $\lambda_0 \geq \frac{(1-\delta)(\delta v+c)}{\delta(1-\delta)(b+\delta V)-\delta^2 v}$. The sufficient condition in the statement of Theorem 1 simplifies this by using the fact that $V \geq \frac{b}{1-\delta}$ and hence provides a condition solely in terms of the primitives. The condition in the theorem says that the relationship constraint (RC) is satisfied when the immediate chance of resolving the clarity problem, λ_0, is high enough relative to the discounted value of the outside option, $\delta v/(1-\delta)$, net of the matching cost c.

3.1 Characterizing behavior when clarity is single-peaked

We now characterize the agent's optimal behavior, imposing Assumption 1 that clarity is single-peaked. Recall that clarity, $(\lambda_t)_{t=0}^{\infty}$, is single-peaked if there exists some period $T \geq 0$ such that $(\lambda)_{t=0}^{T}$ is increasing and $(\lambda_t)_{t=T}^{\infty}$ is decreasing.

Theorem 2. *The sequence of continuation values, $(W_t)_{t=0}^{\infty}$ is single-peaked with peak $\tau \leq T$.*

Theorem 2 shows that continuation values $(W_t)_{t=0}^{n}$ are also single-peaked. Their peak, however, is (weakly) before the peak of $(\lambda)_{t=0}^{n}$. The reason is option value: if λ_T is not much larger than λ_{T-1}, we could have that $W_{T-1} > W_T$ since at time $T-1$ the agent has two actions to try which are likely to be productive, but in period T only one of these is left.

This result is useful in further characterizing the solution to the agent's problem. We now show that there is a deadline, $K < \infty$. Up to period K the agent will search for a productive action with the current principal if shocks are low enough. But if no productive action is found

[17] We cannot have $F(s^*) = 0$, since this would violate constraint (RC).

in period $K \geq T$ the agent immediately terminates the relationship and rematches to a new principal. Thus any relationship which survives period K must be in its productive phase. This is formally stated in the theorem below.

Theorem 3. *There exists a $K \geq T$ such that $\delta W_t \geq \delta W_0 - c$ for all $t \leq K$ and $\delta W_t < \delta W_0 - c$ for all $t > K$.*

Theorem 3, when combined with equation (1), implies a simpler functional form for continuation values

$$W_t(s) = \begin{cases} \max\{s + \delta W_0 - c, \lambda_t (b + \delta V) + (1 - \lambda_t)\,\delta W_{t+1}\} & \text{if } t < K \\ \max\{s + \delta W_0 - c, \lambda_t (b + \delta V) + (1 - \lambda_t)(\delta W_0 - c)\} & \text{if } t = K \end{cases}. \tag{5}$$

It also allows us to obtain a simpler characterization of the cutoff shock

$$s_t^* = \begin{cases} \lambda_t (b + \delta V) + (1 - \lambda_t)\,\delta W_{t+1} - \delta W_0 + c & \text{if } t < K \\ \lambda_K (b + \delta V - \delta W_0 + c) & \text{if } t = K \end{cases}. \tag{6}$$

Since period K is the last period before the agent restarts with a new principal we do not need to specify W_t and s_t^* for $t > K$.

These results explain how the cutoff shocks to the outside option vary across time. In the search phase of a relationship, for periods $t < \tau$, s_t^* increases since both λ_t and W_t are increasing. For periods $t > T$ both λ_t and W_t are decreasing and the cutoff shock decreases, as clarity and the option value of continuing with the current principal are exhausted. In the productive phase, s^* does not depend on t, so relationships are terminated at a constant rate, with $s^* \geq s_t^*$. This is summarized in the following empirical implication.

Empirical Implication 1. *Assuming $\lambda_t < 1$ for all t, in early periods ($t \leq K$) relationships end with a higher probability, which varies with t. In later periods ($t > K$) only productive relationships remain and terminate at a constant rate.*

3.2 Benchmark without clarity problems

As a benchmark, consider the case where the principal and agent do not face clarity problems. This is a special case of our framework where $\lambda_0 = 1$, so the agent immediately knows a productive action for every principal. Then $W_0 = V$ and termination of relationships is governed by the fixed cutoff $s^* = b + \delta V - \delta W_0 + c = b + c$. The probability of a shock above s^* is $1 - F(s^*)$ and is constant because shocks are i.i.d. This is summarized in the proposition below.

Proposition 1. *If $\lambda_0 = 1$, the probability that a relationship ends in any period is constant.*

3.3 Clarity and relationship dynamics

We now study how clarity shapes relationship formation and termination. We first show that the agent's continuation values are increasing in clarity, which implies that agents with higher clarity are more likely to satisfy (RC). This means that such agents are more likely to pay cost $c \geq 0$ to form a relationship in the first place. We then consider how clarity impacts relationship termination by looking at its effect on the cutoff shock s_t^*, the threshold above which agents take the outside option. Lastly, since exporters in our data differ with respect to discount factor, we show that more patient agents are less likely to end productive relationships.

Proposition 2. *$W_{t'}$ is increasing in λ_t for any $t', t \leq K$. Furthermore, W_0 is strictly increasing in λ_t if and only if $F(s_k^*) > 0$, for all $0 \leq k \leq t$ and $\lambda_k < 1$ for all $0 < k < t - 1$.*

The proof of the proposition is in Appendix B.4. A higher λ_t raises the value of both staying in the current relationship (because the agent is more likely to resolve the clarity problem) and rematching (because W_0 is the ex-ante value of starting with a new principal). Strict monotonicity requires that the relationship can stay in the search phase until period t, so that λ_t can matter. This will not happen if $\lambda_k = 1$ for some $k < t$, since the relationship becomes productive in period k with probability 1. Similarly, if $F(s_k^*) = 0$ for some $k < t$, the agent always takes the outside option in period k, and the relationship always ends before period t.

Proposition 2 showed that W_0 is increasing in λ_t and therefore in clarity. Inequality (RC) is more likely to hold when W_0 is large and hence higher clarity makes it more likely that an agent will match with a principal and attempt to form a productive relationship instead of always taking the outside option. This is recorded in the empirical implication below.

Empirical Implication 2. *An agent with higher clarity is more likely to attempt to form a productive relationship.*

Holding clarity fixed, but increasing W_0, the expected payoff from starting a new relationship, makes it more likely that an existing relationship ends. This is true in both the search phase (where the cutoffs s_t^* get smaller) and in the productive phase (where s^* decreases).

Proposition 3. *For each $t \leq K$, the total derivative $\frac{\mathrm{d}s_t^*}{\mathrm{d}W_0} < 0$. Furthermore $\frac{\mathrm{d}s^*}{\mathrm{d}W_0} < 0$.*

The proof is in Appendix B.5. The probability that an agent takes the outside option in the search phase is $1 - F(s_t^*)$. This probability is decreasing in s_t^* and since s_t^* decreases with W_0, the probability that the agent ends a relationship is increasing in W_0. The same holds in the productive phase.

Combining Propositions 2 and 3 yields a stark prediction for productive relationships: higher clarity (namely increasing any component of the sequence $(\lambda_t)_{t=0}^K$) reduces the cutoff s^* for terminating a relationship in the productive phase. This increases the probability that a relationship in the productive phase ends. The next Proposition considers how clarity affects relationship termination during the search phase.

Proposition 4. *For each $t \leq K$, $\frac{\mathrm{d}s_\tau^*}{\mathrm{d}\lambda_t} \geq 0$ if $\tau \leq t$ and $\frac{\mathrm{d}s_\tau^*}{\mathrm{d}\lambda_t} \leq 0$ if $\tau > t$.*

Raising λ_t has two effects during the search phase. As we already saw, it increases the value of starting a new relationship, which makes the agent more likely to terminate the current relationship. However, before period t, it also strengthens the case for staying in the current relationship because the agent can directly benefit from the higher probability that an upcoming attempt is more likely to find a productive action. The key part of the proof, which is in Appendix B.6, is to bound these countervailing effects and to show that the latter dominates before period t. Proposition 4 shows that if some λ_t is increased, the agent is less likely to terminate the relationship in all periods up to (and including) period t, but is more likely to do so after period t. As shown above, this increase in the likelihood of termination continues to hold even if the relationship becomes productive. The next empirical implication records this reversal in termination rates.

Empirical Implication 3. *Agents with higher clarity are less likely to end relationships early, but more likely to end them later, especially after period K when only productive relationships remain.*

To compare this result to Jovanovic (1979), we can interpret an agent with higher clarity in our model as one who draws from a better distribution of worker-firm match productivities in his. Since this comparative static is not derived in Jovanovic (1979), we derive it in Online Appendix B, using a two-period version of his model.[18] In Jovanovic's setting, no reversal in termination rates arises: agents with better match-productivity distributions are less likely to end relationships in all time periods.

Finally, we show that patient agents are less likely to end productive relationships. The key step is showing that the agent's continuation value is increasing in the discount factor.

Proposition 5. *The agent's W_t is increasing in δ, i.e., $\frac{\mathrm{d}W_t}{\mathrm{d}\delta} > 0$ for all $t \le K$.*

The proof of Proposition 5 is in Appendix B.7. Although the proposition shows that the agent's continuation value increases with the discount factor, it does not imply how the probability of terminating a productive relationship varies with δ. In problems of learning such as ours, Quah and Strulovici (2013) show that patient agents may stop experimenting earlier in some cases. However, this counterintuitive result does not obtain in our model. As we show in Appendix B.8, patient agents are in fact less likely to terminate productive relationships.

Empirical Implication 4. *As the agent becomes more patient, he is less likely to end a productive relationship.*

3.4 Extensions

In this section we provide two brief extensions. The first introduces different types of principals. In the empirical analysis, this will allow us to ask whether some agents have lower clarity because they are of a low type or because they interact with low type buyers. We also describe how principal incentives can be added to the model.

3.4.1 Different Principal Types and Components of Clarity

Our notion of clarity captures both the principal's ability to explain what she requires and the agent's ability to understand and deliver it. So far we have assumed that all principals are alike. Suppose now there are different types of principals, denoted by $\theta \in \Theta$. To avoid technicalities, assume Θ is finite. Once the agent is matched with a principal, he observes the principal's type θ and updates his belief about which actions the principal may find productive.

For each principal type, we have a sequence $(\lambda_t^\theta)_{t=0}^\infty$ of probabilities that the agent finds a productive action in period t, conditional on not having found one before. This sequence

[18] The two-period version simplifies the derivation, though the same result obtains in the infinite-horizon case.

represents the clarity problem when the agent is matched to a principal of type θ. As before, all we need to assume is that $(\lambda_t^\theta)_{t=0}^\infty$ is single-peaked. However, we now provide an additional microfoundation which will be useful for empirics.

When a principal of type θ matches with the agent, there are $|\mathcal{A}_i| = N$ possible actions. Suppose that all actions could be productive if performed correctly. However, the principal may not know how to perform the action correctly or may be unable to explain it clearly. Furthermore, the agent may not be able to understand what is required.

The agent's effective ability to understand what the principal wants is a random variable, $\theta_a + \varepsilon_a^k$, $\theta_a \in \mathbb{R}$. The net difficulty of understanding action k is $-\theta_p^k + \varepsilon_p^k$, where ε_a^k and ε_p^k are i.i.d. *Gumbel* $(0, 1)$. We interpret $\theta_p^k \in \mathbb{R}$ as the action's overall simplicity. This parameter captures both the intrinsic complexity of action k and the principal's ability to identify that the action is productive and effectively communicate how to perform it. Action k is correctly performed, and hence productive, if the agent's ability to understand exceeds the difficulty of understanding the action. This occurs with probability

$$\lambda_k^\theta = \Pr\left[\theta_a + \varepsilon_a^k > -\theta_p^k + \varepsilon_p^k\right] = \frac{1}{1 + e^{-\left(\theta_a + \theta_p^k\right)}}, \tag{7}$$

where the last equality holds since both ε_a^k and ε_p^k are i.i.d. Gumbel random variables, so their difference is logistic. Note that λ_k^θ varies across different actions k, as well as principal-agent pairs. If the probability that each action is productive is independent of other actions, we will get a decreasing (and thus single-peaked) sequence of $(\lambda_k^\theta)_{k=0}^\infty$. This provides a simple logit microfoundation for clarity which divides it into an agent component, a principal component, and idiosyncratic noise. In Section 5.2 we use exactly this decomposition in our estimation.

Our previous results go through as long as we are careful to condition on θ where appropriate. In particular, the cutoff for taking the outside option in period t, $s_t^{*\theta}$, will depend on θ, as will K^θ and the value functions W_t^θ. Note however, that W_0 will now equal $\mathbb{E}[W_0^\theta]$ and not simply W_0^θ. This reflects the fact that when the agent is matched to a new principal, a new θ is drawn. Observe that V will not depend on θ, since by definition, the agent knows a productive action for the principal at this point and only the expected value of W_0 matters.

The model is solved in a similar way: start with a guess for $W_0 = \mathbb{E}[W_0^\theta]$, then for each θ compute the associated W_t^θ as before and finally take expectations over θ to get an updated W_0. Theorem 1 extends to this case, by taking expectations over principal types. The comparative statics also extend to this environment.

Allowing for different principal and agent types yields more realistic predictions about when relationships end. For example, not all relationships that fail to find a productive action will end, on average, on the same shipment. The last period that an unproductive relationship can survive (K_θ) will depend on clarity,$(\lambda_t^\theta)_{t=0}^\infty$, which is a function of the principal's and agent's type. When averaging across different types in our data, if $\Pr(\theta : K_\theta = t)$ is decreasing in t, we would expect the probability of a relationship ending to be the highest in the first shipment and then to decrease over time. Our comparative statics on clarity, however, still hold. The one most relevant for our empirical application is that agents with higher clarity terminate productive relationships more often. This is a straightforward extension of Proposition 4 and can be tested by conditioning on relationships that survive past period $\max_\theta K_\theta$, when only productive relationships remain.

3.4.2 Principal Incentives

By assuming that the principal is committed to a truthful strategy, which pays the bonus if and only if the agent chooses a productive action, we isolated the agent's problem. However, the model can be extended to allow the principal to have a more active role. One approach is to embed our model in a competitive markets environment, where principals (firms) and pay agents their marginal productivity (e.g., Jovanovic, 1979; Bonatti and Hörner, 2017). Another possibility is to introduce a standard moral hazard environment, where contracts can depend on the output produced but not on the agent's outside option. In this case, the principal-optimal contract will pay the agent a fixed bonus if and only if a productive action was taken.

Alternatively, we can embed our model into a relational contracting framework (e.g., Board, 2011). The preceding analysis characterizes an agent self-enforcing contract and so we need to add dynamic principal incentives. We take this route now and also consider the principal's strategic response to our main comparative static—the fact that higher clarity agents are more likely to abandon a productive relationship later on.

To simplify the exposition suppose that $\lambda_0 > 0$, but $\lambda_k = 0$ for $k \geq 1$. So the agent has at most one action to try and if it is unproductive the relationship ends. When in a relationship, the principal gets utility $\varsigma > 0$ in period t if the agent took a productive action, and zero otherwise. If the relationship ends, the principal gets some exogenous continuation value $\rho > 0$.[19] We restrict attention to Markov perfect equilibrium (MPE), where the state is the agent's action set and beliefs over which actions are productive for the current principal, and are interested in characterizing the agent-optimal MPE.

Every efficient MPE maximizes the value of the current relationship, since this also maximizes continuation values across the entire game.[20] For this reason, transfers to the agent will only occur after a productive action is chosen, since this minimizes the agent's incentive to terminate the relationship. Let $R(\lambda_0)$ be the principal's continuation value if she is matched to an agent with clarity λ_0. If the principal pays bonus $b \in \mathbb{R}$ in every period a productive action was chosen, her continuation value from a productive relationship is $R_p(\lambda_0) = F(s^*)(\varsigma - b + \delta R_p(\lambda_0)) + (1 - F(s^*))\delta\rho$. The principal is willing to pay bonus b to the agent instead of keeping surplus and being unmatched, as long as $\varsigma - b + \delta R_p(\lambda_0) \geq \varsigma + \delta\rho$. In the agent-preferred MPE this transfer is as large as possible, so that $b = \delta R_p(\lambda_0) - \delta\rho = \delta F(s^*)\varsigma - \delta(1-\delta)\rho$; the principal pays the agent this b in every period when a productive action is taken and nothing otherwise. Observe that if the agent never terminates a productive relationship b only depends on exogenous parameters. Since only $\lambda_0 > 0$, the principal faces the same incentive problem the first time the agent chooses a productive action as when the players are already in a productive relationship. Thus, the above choice of b ensures that the principal is also willing to reveal

[19]This could be from matching to a different type of agent (e.g., if the clarity distribution is absolutely continuous) or, in our applications, from the principal either grading exams herself or buying roses at auction.

[20]For certain parameters, this will also happen in every subgame perfect equilibrium since clarity problems diminish the value of a new relationship relative to an existing one. See McAdams (2011) for a discussion of the difference between overall welfare-maximization and welfare-maximization within a relationship.

that a productive action was found during the search phase. Online Appendix A shows how to determine s^* as well as the other key variables in this extension of the model.

An implication of Propositions 2 and 3 is that higher clarity agents end productive relationships more often, i.e., s^* is decreasing in λ_0. However, an agent with a higher λ_0 is also more likely to reach a productive relationship in the first place (and there could also be secondary effects through the impact of s^* on b), so the effect on the principal's profit is unclear. Online Appendix A shows that our model can generate a non-monotonicity in $R(\lambda_0)$: for small λ_0 the principal's continuation value is increasing, but for large λ_0 it is decreasing. This means that if the principal incurs a fixed cost (e.g., training) after an agent matches to her but before work can begin, she may decide not to engage with high clarity agents.

The fact that employers are wary of high clarity (or high ability) agents is well documented. For instance, popular commentary reflects these concerns: "The truth is, when hiring managers say, 'You're overqualified,' what they really mean is something else. Maybe they think you'll leave as soon as something better comes along" (Forbes, 2025). The U.S. Court of Appeals upheld that a police department had shown a "rational basis" to disqualify an applicant for a position due to over-qualification (Jordan v. City of New London, 2000). In effect, the court deemed it lawful for the employer to discriminate against high-type individuals because they are more likely to leave the job. In our empirical application, we cannot test this strategic response by buyers because we lack the full set of transactions, and we do not observe contacts that do not lead to an initial shipment.

4 Context

4.1 Data

Our empirical analysis uses transaction-level data on Ethiopian flower exports and two exporter censuses. We now describe these data sources, our sample construction and key variables.

Transaction-level Data. Our main data source is the standard ASYCUDA customs data of Ethiopian flower exports between 2007-2019.[21] This data includes the transaction date, exporter, buyer, commercial description of the product (e.g., roses), net weight, free-on-board (FOB) value in USD, FOB value in Ethiopian Birr, and quantity in stems, among others.

The data includes deanonymized exporter names compiled manually, but they are subject to typos and variants of the same exporter's name. To standardize names and classify each firm as foreign- or domestically-owned, we triangulate several independent sources.[22] We start with the Ministry of Trade & Industry (MoTI) company registration database, which lists registered owners, and validate classifications using firm listings and staff interviews at the Ethiopian Horticulture Development Authority (EHDA), the Ethiopian Horticulture Producer Association (EHPA), and the Ethiopian Investment Agency (EIA).

[21] In the realm of customs and trade, the Automated System for Customs Data (ASYCUDA) is a well-established system developed by the United Nations Conference on Trade and Development (UNCTAD).

[22] Less than 5% of exporters were joint ventures, which we designate as foreign-owned, because, in our interviews, we learned that the foreign partner is typically in charge of direct sales and marketing in export markets.

Each exporter-buyer pair that transacts is considered a relationship, and we call their first shipment an *attempt*. Relationships have two phases. In the first phase, the exporter searches for a productive action for the buyer. If this is successful, they enter into the productive phase. In Section 5, we show that relationships surviving the fourth shipment are most likely in the productive phase. As such, we label a relationship as productive after the fourth shipment.[23] We consider a relationship to have ended when the pair stops transacting for at least six months.[24] In our analysis of relationship termination, we include only relationships that could potentially end within our data, i.e., those that start at least six months before our sample ends.

While either party can end a relationship, we attribute terminations to the exporter, since interviews with exporters suggest they initiate most terminations. Moreover, we study terminations when auction prices are high, which is when exporters benefit from ending the relationship at the buyer's expense. Because relationship prices are usually fixed throughout the year, a higher auction price increases the exporter's incentive to sell at the auction and raises the buyer's cost of sourcing flowers outside the relationship. Thus, it is plausible to attribute the rise in relationship terminations to exporters during periods when exporters face a stronger incentive to defect, while buyers have a weaker incentive to do so.

Sample Construction. The customs data includes 274,065 flower export transactions from July 2007 to July 2019. We first restrict the sample to roses, the dominant flower crop, leaving us with 208,639 transactions. Next we drop a further 7,076 transactions (or 3.4%) of transactions that (i) lack a buyer, (ii) report implausible unit prices or weights, (iii) involve logistics agents, unpackers, or flower traders, who export but do not produce flowers. We exclude shipments to foreign-owned exporters' headquarters, as they are neither relational sales nor auctions (56,766 transactions). To observe the entire life of each exporter–buyer pair, we exclude relationships that begin before July 2008 (15,826 transactions) and after July 2018 (8,034 transactions).[25] Our final sample consists of 120,937 transactions across 75 exporters.

Exporter Censuses. We also draw on firm-level censuses administered to key management staff in 2007/08 and 2010/11 (Mano et al., 2011). They provide insights into exporters' production capabilities and processes, including varieties produced, yields, arable land, altitude, greenhouse size, and quality-control measures. The censuses covered 64 exporters in 2008 and 77 in 2011.[26] They were conducted jointly by the National Graduate Institute for Policy Studies (GRIPS), the Foundation for Advanced Studies on International Development (FASID), and the Ethiopian Development Research Institute (EDRI), with the aim of understanding Ethiopia's rapidly expanding flower sector and its potential for poverty reduction.

[23] Our empirical results are robust to using the third or fifth shipment as well.

[24] Field interviews with exporters indicate that if no transactions occur for six-months, the relationship has most likely ended. Our results are robust to using alternative time frames. We define relationship termination in this way, because the absence of future transactions means relationships that ended earlier have longer observable periods of inactivity. Most relationships transact semiweekly, weekly or fortnightly.

[25] We need to exclude relationships that start too early in the sample because their first *observed* transaction may not necessarily be their first actual transaction. We also need to exclude relationships that start too late in the sample because we need six months without a transaction to identify the exporter ended the relationship.

[26] There were 67 flower farms in operation in 2008 and 85 in 2011.

4.2 Industry Background

This section reviews Ethiopia's cut flower industry, its contractual practices, relationship patterns, and similarities and differences between foreign-owned and domestically owned exporters.

Flower Industry. Ethiopia entered the cut-flower market in 2004 and became the fifth-largest exporter worldwide, with annual exports often exceeding USD 250 million.[27] Cut-flower sales are a key source of foreign exchange and jobs, contributing 8-14% of the country's merchandise-export earnings and directly employing roughly 40,000 workers (Oqubay, 2015).

Sales channels. The Ethiopian flower season runs from mid-August to mid-August of the following year. Exporters reach buyers via two channels: (i) auctions, chiefly in the Netherlands, and (ii) direct sales that rely on relationships with wholesalers. Logistics and phytosanitary rules are similar, but institutional support differs. The auction grades every lot, guarantees payment and delivery, and has no volume commitments. Auction sales incur intermediary costs, such as handling fees and additional transport costs to and from the Netherlands. Direct relationships avoid these fees but involve counterparty risk.[28] In relationships, exporters and buyers sketch a custom plan at the start of the season, including monthly volume, flower varieties, prices and payment terms. While these agreements allow some flexibility, prices are typically fixed.

Relationships are desirable for exporters because they help avoid auction-related costs, offer higher average prices, and reduce price uncertainty. Appendix Figure D.1 shows that average relationship prices are generally higher than auction prices.[29] Indeed, the number of relationships and their associated sales have grown over time, largely because foreign-owned exporters have shifted away from auctions (Figure 1). Domestically owned firms were less successful at developing new relationships. While both groups send about 170 shipments annually, foreign-owned exporters send 70% of their shipments directly to buyers, compared to only 21% for domestically-owned firms (Appendix Table C.1). Foreign-owned exporters also attempt twice as many new relationships and maintain twice as many ongoing ones.

Reaching Productive Relationships. Auction sales must comply with the Dutch Flower Auctions Association's standards for quality, grade, stem length, packaging, and product coding. The rules and standards are stable. Unlike auction sales, where exporters need to learn these rules only once, *each* direct relationship involves distinct buyer-specific requirements.

Buyer interviews indicate that, while they have an idea of their desired product, translating that vision into actionable instructions for exporters is far from straightforward, creating frequent opportunities for misunderstanding and misalignment in execution. Exporter interviews reveal that buyers request ripeness, stem length, and post-harvest treatment tailored to their downstream logistics. Packaging is also customized: buyers specify bunching, assortments, varietal, flower food and preservatives, sleeves and wraps to create a particular in-transit

[27] National Bank of Ethiopia, Annual Reports: https://nbe.gov.et, accessed May 2025.

[28] Cross-border transactions invite opportunism: exporters may ship sub-standard flowers and buyers may refuse payment. External enforcement of contracts is limited due to unclear jurisdiction and weak legal mechanisms.

[29] Since relationship prices were not consistently higher prior to 2012, we conduct a robustness check restricting the sample to shipments from 2012 onward. The core findings of the paper remain unchanged.

Figure 1: Sales Channels in the Ethiopian Floriculture Industry

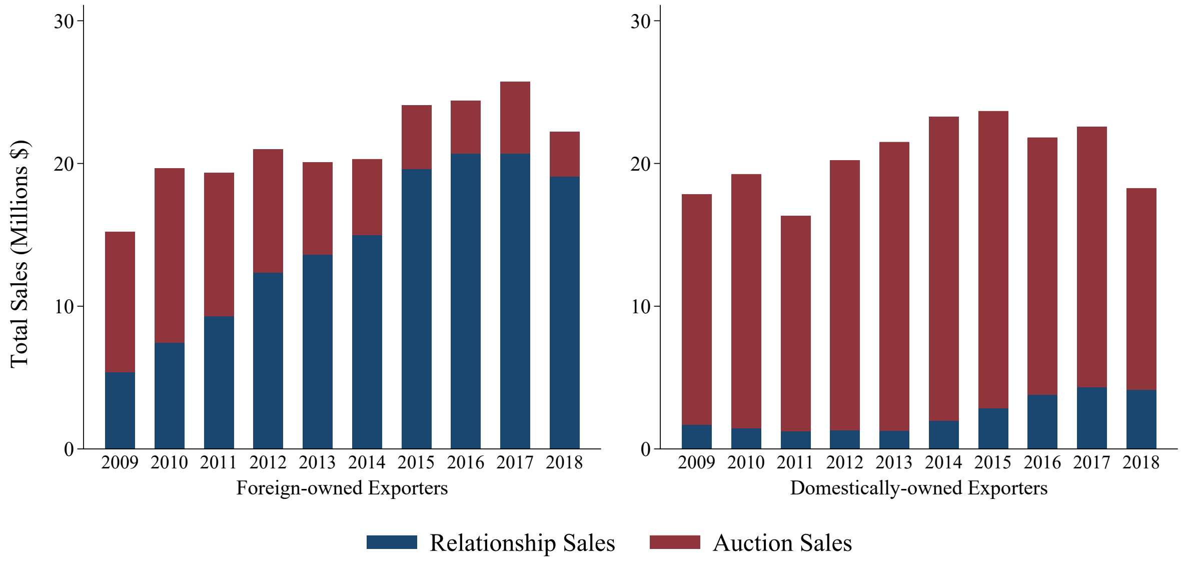

Note: The figures display seasonal sales, dividing them between direct transactions (relationships) and auctions, separately for foreign- (left side) and domestically-owned exporters (right side). Each season runs from mid-August to mid-August of the following year, except for the 2018 season, which spans only from mid-August to June 30, 2019, due to data availability.

micro-climate (temperature and humidity). Some buyers favor "compact packaging" — tight bundles in cardboard boxes — whereas others reject it for fear of "overfilling" and bruising.

Scale presents an additional hurdle. Demonstrating competence on a trial shipment does not guarantee reliable performance at higher volumes. As a result, the "starting small" strategy stressed elsewhere (e.g., Ghosh and Ray, 1996; Kranton, 1996; Watson, 1999) is less applicable in this context. In our data, the first three shipments are only 3 to 5% smaller than later ones (Appendix Table C.2).

Consistent with widespread clarity problems, approximately half of all relationships fail before the fourth shipment (Figure 2). This issue is more pronounced among domestically owned exporters, 50% of whose relationships terminate after the first shipment – compared to 30% for foreign-owned exporters. However, once a relationship reaches the fourth shipment, retention rates are similar across the two types of firm ownership.

Domestic and Foreign Exporters. Foreign- and domestically owned rose exporters in Ethiopia are similar in production capabilities (scale of operations, input use, and product quality), but they differ in key attributes that affect relationship building and termination. For example, their discount factors and their experience dealing with international buyers. We begin by outlining similarities, then turn to key differences.

Production scale (including direct sales and auctions) across ownership types shows a balanced distribution (Appendix Figure D.2). Of the four largest firms, two are foreign-owned; among the top ten largest, four are foreign-owned. Two independent firm censuses confirm that cultivated hectares and yields are statistically indistinguishable (Appendix Table C.3).

Rose quality does not differ by firm ownership. Quality depends on variety, from low-end Sweetheart to high-end Tea-hybrid. Foreign- and domestically owned exporters allocate similar

Figure 2: Relationship Survival by Shipment Number

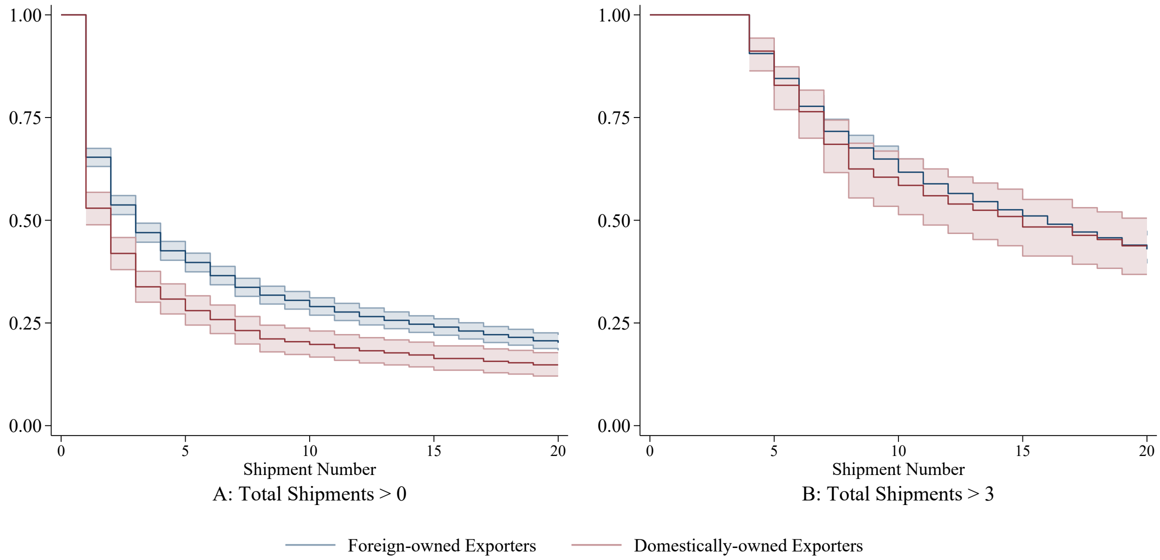

Note: The figures display the estimated Kaplan–Meier survival function for relationships between Ethiopian flower exporters and their buyers, distinguishing whether the exporter is foreign-owned (blue) or domestically-owned (red). A relationship is considered terminated following a shipment if no further transactions occur within the subsequent six months. The left-hand panel includes all direct relationships, while the right-hand panel restricts to relationships that have reached at least the third shipment. Each of the survival graphs is shown with its corresponding 95% pointwise confidence interval.

land shares to each variety (see Appendix Table C.3). Although customs records omit variety, we find no systematic quality gap based on stem weight—a standard proxy for quality in the industry, as heavier heads indicate superior quality and higher market value (Appendix Table C.4).[30] Likewise, we find no differences in farm altitude, greenhouse hectares, imported-fertilizer use, customer rejection rates or quality control practices; however, domestically-owned exporters put more weight on buyer feedback and culling defective stems (see Appendix Table C.5).

Although similar in many ways, domestically owned exporters face significantly greater credit constraints than foreign-owned firms. In our census data, 75% of domestically-owned firms list financing as a major obstacle – nearly twice the rate of foreign-owned exporters. They must post collateral averaging 97% of the loan value, pay approximately one percentage point higher interest rates, rely on external finance to cover over half of their first-year operations (versus 32% for foreign-owned exporters; see Appendix Table C.6). These credit constraints imply a lower discount factor for domestic exporters. These exporters also differ in their experience dealing with international buyers, which could affect clarity. All foreign-owned firms bring prior floriculture experience, some spanning generations, that ranks them among the world's most successful rose sellers. By contrast, only two of 31 domestic entrants had any acquaintance with international buyers, and neither had exported roses.[31]

[30] Macchiavello and Morjaria (2021) also use unit stem weight as a proxy for quality.
[31] Source: author interviews and Oqubay (2015).

5 Empirical Results

We next empirically test our theoretical predictions regarding how exporters form and end relationships. In line with empirical implication 1, we provide evidence that clarity problems are prevalent in the industry. Specifically, we show that relationships are significantly more likely to end during their initial search phase than later on, when only productive relationships remain. Additionally, we find that domestically owned exporters struggle more to reach productive relationships; their relationships are more likely to end within the first three shipments.

We showed that a lower discount factor leads to a higher probability of terminating a productive relationship (empirical implication 4). Taken alone, this suggests that domestically owned exporters, who have lower discount factors, are more likely to end productive relationships in response to increases in the auction price. However, this prediction may not hold in our model because domestically owned exporters also face greater difficulty in reaching the productive phase of relationships due to lower clarity. This difficulty makes them less willing to terminate productive relationships, since it is harder for them to build new ones (empirical implication 3). We find that this clarity effect dominates the discount factor impact. Despite having lower discount factors, domestically owned exporters are less likely to end productive relationships in response to positive shocks to their outside option. Furthermore, they are also less likely to initiate building a relationship in the first place, consistent with the implications of empirical implication 2.

In Section 3.4.1, we developed a microfoundation for clarity that follows a logit distribution as a function of agent and principal types. To bring this to the data, we decompose the likelihood of reaching the productive phase of a relationship into an exporter and a buyer component using the methodology of Abowd et al. (1999), henceforth AKM. Using the estimated exporter and buyer clarity components, we assess whether domestically owned exporters struggle to form productive relationships because of having low clarity, or facing buyers with low clarity, or both. We find that domestically owned exporters do not face systematically different buyer types. Instead, their lower success rate in forming productive relationships is driven by a lower exporter clarity component.

Finally, we test our key comparative static, empirical implication 3: exporters who are better at forming productive relationships are also more likely to exit those relationships. Specifically, we show that, in response to positive shocks to the outside option, exporters with *higher* clarity – that is, those better at forming productive relationships – exhibit a higher likelihood of terminating existing productive relationships.

5.1 Reaching and Ending Productive Relationships

Reaching Productive Relationships. Empirical implication 1 states that when clarity is a problem, relationships are more likely to end before becoming productive, whereas termination occurs at a constant rate after relationships reach their productive phase. We now estimate the probability that a relationship between exporter e and buyer b ends at shipment h relative to the probability of ending at the 10^{th} shipment. This is implemented by the following regression specification

$$\mathbb{1}[Rel.\ End]_{e,b,h} = \sum_{i=1, i\neq 10}^{30} \beta_i \mathbb{1}[h=i]_{e,b,h} + \phi_e + \zeta_t + \epsilon_{e,b,h}. \tag{8}$$

The estimation includes exporter fixed effects (ϕ_e) to control for time-invariant differences across exporters, and year-by-month fixed effects (ζ_t) to absorb common industry shocks. As a result, the β_i's coefficients capture the within-exporter variation in the probability of relationship termination across shipment numbers, relative to the baseline of the 10$^{\text{th}}$ shipment.

Figure 3: Probability of the Relationship Ending Relative to the 10^{th} Shipment

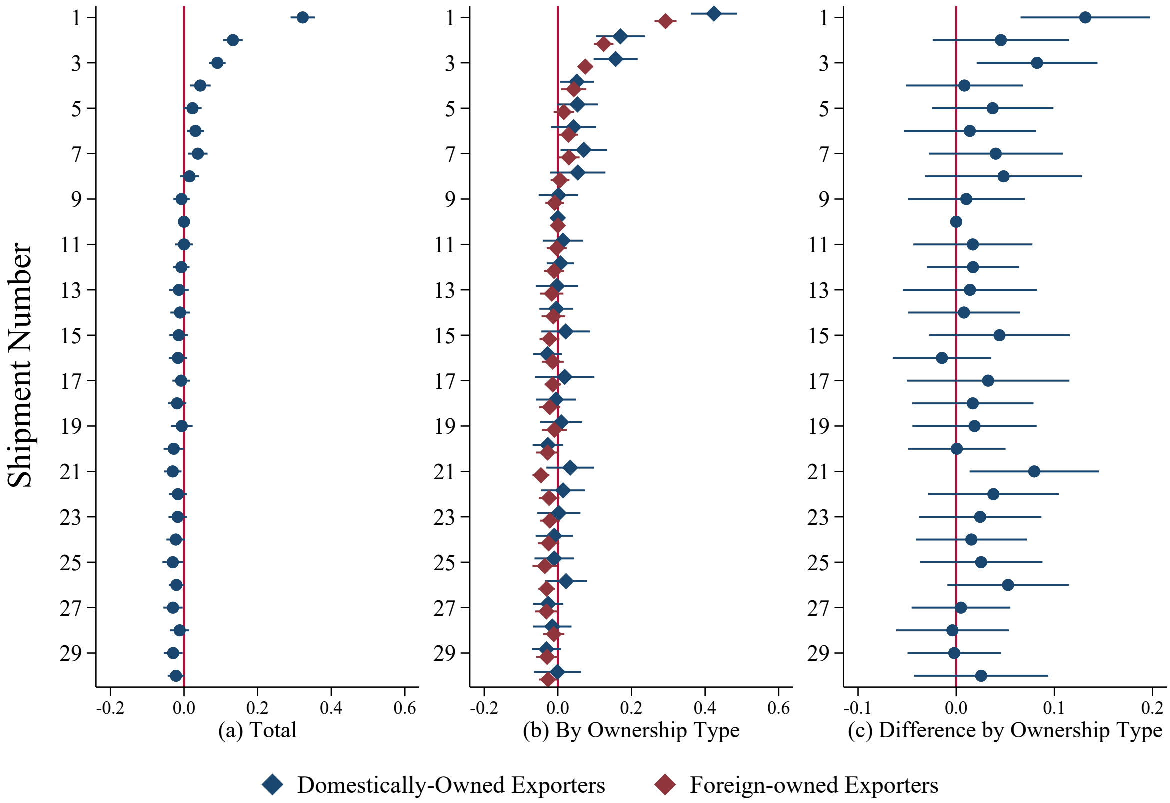

Note: The figure presents the estimated probability that a relationship ends at a given shipment number relative to the 10th shipment. Panel (a) plots the $\hat{\beta}_{1,i}$ coefficients from equation 8, capturing the probability of termination by shipment number. Panel (b) displays results from equation 9, separating estimates for domestically owned exporters ($\hat{\beta}_{1,i}$) and foreign-owned exporters ($\hat{\beta}_{1,i} + \hat{\beta}_{2,i}$). Panel (c) shows the differences in termination probabilities between foreign- and domestically owned exporters, as captured by ($\hat{\beta}_{2,i}$). Standard errors are two-way clustered at the exporter and buyer levels, and all coefficients are displayed with 95% confidence intervals.

We also investigate whether early relationship terminations are more prevalent among domestically owned exporters. To do so, we interact the shipment dummies with an indicator variable, D_e, which equals 1 if exporter e is domestically-owned, as follows

$$\mathbb{1}[Rel.\ End]_{e,b,h} = \sum_{i=1, i\neq 10}^{30} (\beta_{1,i} \mathbb{1}[h=i]_{e,b,h} + \beta_{2,i} \mathbb{1}[h=i]_{e,b,h} \times D_e) + \phi_e + \zeta_t + \nu_{e,b,h}. \tag{9}$$

Figure 3a confirms that relationships are substantially more likely to end during the first 3 or

4 shipments, consistent with exporters facing difficulty in reaching the productive phase. After this point, the probability of ending a relationship remains approximately constant. Quantitatively, the probability of termination in the first four shipments is, on average, 14 percentage points higher relative to the 10^{th} shipment. In contrast, for shipments 5 to 9 and 11 to 30, the average termination likelihood is 0.01 percentage points lower, and mostly insignificant.

Figure 3b shows that domestically-owned exporters experience significantly more early exits than foreign-owned exporters. Figure 3c displays these differences: in the first shipment alone, domestically-owned exporters are 13 pp more likely to end a relationship. Cumulatively, over the first three shipments, they face a 26 pp higher likelihood of relationship termination compared to their foreign-owned counterparts.

Domestic vs Foreign-Owned Exporters: Ending Productive Relationships. Empirical implication 4 shows that exporters with a lower discount factor are more likely to end a productive relationship when there is a positive shock to the outside option. On its own, this implies that domestically owned exporters – who have a lower discount factor – should be more likely to abandon productive relationships in response to positive shocks to the outside option.[32] However, this prediction may not hold in our model because domestically owned exporters also find it more difficult to reach productive relationships (they have lower clarity). In particular, empirical implication 3 shows that such firms are less likely to terminate productive relationships.

To test whether positive shocks to the outside option affect domestically- and foreign-owned exporters differently, we use the price spread between auctions and direct transactions as an exogenous shifter of the value of the outside option. The higher the auction price, the stronger an exporter's temptation to send a shipment to the auction instead of their regular buyer. Importantly, the auction price is exogenous from the perspective of Ethiopian exporters, as they are small relative to the global flower market and cannot influence prices. We estimate

$$Y_{e,t} = \beta_0 + \beta_1 D_e + \beta_2 Price\ Spread_{t+1} + \beta_3 Price\ Spread_{t+1} \times D_e + \Gamma X_{e,t} + \epsilon_{e,t}, \qquad (10)$$

where the dependent variable, $Y_{e,t}$, is either the number of relationships ending, or an indicator for exporter e ending at least one relationship in week t. Dummy variable D_e indicates whether exporter e is domestically owned or not, and $Price\ Spread_{t+1}$ denotes the spread between the average auction price and the average price in direct shipments for the following week.

We use the price spread for the following week ($t + 1$) because if the last shipment of a relationship occurs in week t, it implies that the exporter chose not to ship flowers in week $t+1$. We restrict the sample to productive relationships, defined as those that have reached the fourth shipment, to ensures terminations are not due to a failure to find a productive action. The set of controls $X_{e,t}$ includes the number of active productive relationships in all specifications, since the number of productive relationships an exporter can end is mechanically constrained by how many they currently maintain. We also report results that include additional controls for the weekly quantity and value of flowers sold by the exporter (see Appendix Table C.7).

Table 1 shows that a one standard deviation increase in the average weekly price spread between auctions and direct relationships raises the probability that a foreign-owned exporter

[32] Appendix Table C.6 shows that domestically owned firms have a lower discount factor.

Table 1: Domestically Owned Exporters and Ending Productive Relationships

Dependent Variable:	Number of Relationships Ending ≥ 1 (Dummy)		Number of Relationships Ending (Count)			
	(1)	(2)	(3)	(4)	(5)	(6)
Price Spread (Std)	0.009**	0.011**	0.011**	0.012**	0.150	0.140*
	(0.004)	(0.004)	(0.005)	(0.005)	(0.092)	(0.075)
Price Spread (Std) x	-0.009*	-0.009*	-0.010*	-0.010	-0.402**	-0.326*
I[Domestic Exp.]	(0.005)	(0.005)	(0.006)	(0.006)	(0.157)	(0.177)
Mean Dep. Var	0.047	0.047	0.053	0.053	0.053	0.059
Observations	17,489	17,488	17,489	17,488	17,489	15,611
Exporter FE		Y		Y		Y
Estimation	OLS	OLS	OLS	OLS	PPML	PPML

Note: The table reports the OLS and PPML estimates of β_2 and β_3 at the exporter-week level from equation 10. In Columns 1-2, the dependent variable is a dummy that equals one if the exporter had at least one relationship end during the week. In Columns 3-6, the dependent variable is the number of relationships that ended that week. Columns 5-6 are estimated using a Poisson Pseudo-Maximum Likelihood (PPML) estimator. A relationship is considered ended if no further shipments are observed between an exporter and a buyer for at least six months. *Price Spread* is defined as the standardized difference between the average auction price and the average price in direct shipments for the following week. The variable *Domestic Exp.* is a dummy indicating whether the exporter is domestically owned or not. The sample includes productive relationships – relationships that end in the first three shipments are not considered. All specifications include a control for the number of active, productive relationships the exporter has that week. Standard errors (in parentheses) are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at a 0.10, 0.05, and 0.01 level, respectively.

ends at least one relationship by 0.011 pp (a 23% increase). Estimates of the effect on the number of relationships ending are consistent across both linear and Poisson models. In the linear specification, the number of relationships ending increases by 0.012 (a 23% increase), while the Poisson specification yields a 15% increase.

Despite a lower discount factor, domestically owned exporters are less likely to terminate relationships in response to positive shocks to the outside option. The effects observed for foreign-owned exporters largely dissipate among domestically owned ones. Empirically, the clarity problems these firms face dominate the effect of their lower discount factor.

It is important to highlight two limitations of the current analysis. First, we classify exporters as either foreign or domestically owned, based on the tendency that foreign-owned exporters are better at building relationships, on average. However, this ignores heterogeneity across exporters within each firm type: some domestically owned exporters are successful in relationship-building, while some foreign-owned firms struggle. Foreign-owned exporters, on average, sell 66% of their production through direct relationships, compared to only 19% for domestically owned exporters. Yet, the top eight domestically owned exporters – ranked by share of direct sales – achieve an average of 65%, nearly matching the foreign-owned average. Conversely, 27% of foreign-owned exporters fall below the average for domestically owned firms.

Second, domestically owned exporters may find it harder to reach productive relationships either because they are less able to meet buyers' needs or because they tend to transact with buyers who are harder to satisfy. That is, the composition of buyers served by domestically versus foreign-owned exporters may differ in ways that contribute to disparities in the likelihood of reaching productive relationships.

The following sections refine the analysis by addressing these limitations. First, we decompose clarity into exporter-specific and buyer-specific components, providing measures of the likelihood of forming productive relationships at both the exporter and buyer levels. We then use these measures to determine whether domestically owned exporters struggle to establish productive relationships due to their own lower ability or because they are matched with worse buyers. Moreover, we use the measure at the exporter level to test whether exporters who excel at forming productive relationships are also more likely to abandon them.

5.2 Decomposing the Likelihood of Reaching Productive Relationships

We now decompose *clarity* into an exporter-specific and buyer-specific components, following the logit microfoundation introduced in Section 3.4.1. Starting from equation 7 and replacing the agent and principal subscripts with e and b (denoting exporter and buyer, respectively), clarity is modeled as a function of the exporter and buyer components

$$\lambda(\theta_e, \theta_b) = \frac{1}{1 + e^{-(\theta_e + \theta_b)}}.$$

The exporter component, θ_e, captures the exporter's inherent ability to deliver productive actions to buyers. This ability may vary due to differences in managerial practices, capability to understand buyer requirements, experience with international buyers, and communications skills, among other factors. The buyer component, θ_b, reflects the buyer's flexibility: a higher θ_b implies that it is easier for any exporter to find a productive action with that buyer. Variation in θ_b may stem from the adaptability of the buyer's product mix (e.g., their ability to sell a broader range of product variants) or logistical and warehousing flexibility.

First, we derive a linear relationship by dividing both sides of the equation by $1 - \lambda$ and taking logs, yielding $Z = \log[\lambda/(1 - \lambda)] = \theta_e + \theta_b$. We then replace each term with its empirical counterpart. The exporter and buyer components (θ_e and θ_b) are the exporter and buyer fixed effects. We also introduce an error term, $\varepsilon_{e,b}$, to capture the idiosyncratic component of the exporter-buyer match. Because λ and thus Z are not directly observable, we substitute Z with an empirical proxy: a binary indicator for whether an exporter-buyer pair reaches the fourth shipment, which is highly correlated with reaching the productive phase of a relationship.[33] The resulting estimating equation is

$$\mathbb{1}[Productive]_{e,b} = \theta_e + \theta_b + \varepsilon_{e,b}. \tag{11}$$

This methodology mirrors the two-way fixed effects framework introduced by Abowd et al. (1999) (AKM) to decompose wages into worker and firm components. The approach has since been widely adopted in the labor economics literature (e.g., Song et al., 2019; Card et al., 2013). We estimate the model using the leave-one-out connected set, which includes only exporters and buyers that remain connected after the removal of any exporter or buyer. We do so to avoid that the coefficients and variances in the following analysis suffer from a bias due to a linear combination of unobserved match-specific variances (Kline et al., 2020).

[33] We define a productive relationship as one that reaches the fourth shipment, consistent with the evidence we presented in Figure 3. Our results are robust to alternative definitions of productive relationships, including the third, fifth, or sixth shipment.

After obtaining estimates of the exporter and buyer components, we examine their role in predicting the likelihood of reaching a productive relationship.[34] We estimate

$$\mathbb{1}[Productive]_{e,b} = \beta_0 + \beta_1 \hat{\theta}_e + \beta_2 \hat{\theta}_b + \epsilon_{e,b}. \tag{12}$$

Both the exporter and buyer components significantly contribute to reaching a productive relationship (Appendix Table C.8). This result holds across the four specifications, which differ in the minimum number of shipments required to define a productive relationship. Our preferred specification, in Column 2, uses a threshold of at least four shipments, based on the evidence in Figure 3, where relationships seem to enter their productive phase by the fourth shipment. In this specification, a one-standard-deviation increase in the exporter component raises the probability of reaching a productive relationship by 21 pp. Similarly, a one-standard-deviation increase in the buyer component leads to a 29 pp increase.

Furthermore, the analysis of variance (ANOVA) reported in Appendix Table C.8 reveals that the buyer component explains approximately 32% of the variance in the likelihood of a relationship becoming productive, while the exporter component accounts for 16%. These results highlight the importance of both sides of the relationship, with the buyer component having greater explanatory power in determining the likelihood of success.

Domestically Owned Firms and the Exporter and Buyer Components. Domestically owned exporters may struggle to reach productive relationships due to their lower clarity (θ_e), exposure to lower-clarity buyers (θ_b), or both. To test for systematic differences in these components between domestically and foreign-owned exporters, we estimate the following equation at the exporter-buyer pair level

$$\hat{\theta}_j = \beta_0 + \beta_1 Domestic_e + \nu_{e,b} \qquad \text{for } j \in \{e, b\}. \tag{13}$$

Our estimates indicate that the primary driver of the difference in the likelihood of reaching a productive relationship between foreign- and domestically owned exporters is the exporter component: domestic firms have significantly lower θ_e, indicating they are less effective at delivering productive actions to buyers. Table 2 presents estimates of equation 13 using four alternative definitions of reaching a productive relationship.[35] The results are consistent across all specifications: domestically owned exporters exhibit exporter component values between 0.5 and 0.85 standard deviations lower than foreign-owned firms, while there is no statistically significant difference in the clarity of the buyers they face.

While the earlier ANOVA results emphasized the importance of the buyer component in reaching productive relationships, the evidence here suggests that domestically owned exporters have lower exporter clarity but do not systematically face lower type buyers. Thus, controlling for the buyer should decrease the severity of the clarity problem, but the difference in clarity between domestically and foreign-owned exporters should persist. To test these predictions, we re-estimate equation 8, now including buyer fixed effects to control for the buyer component.

[34] We winsorize both components at the 5% level to remove extreme values that may bias further analysis.

[35] These definitions correspond to relationships that reach the 3^{rd}, 4^{th}, 5^{th}, or 6^{th} shipment.

Table 2: Clarity Components and Domestically-Owned Exporters

Dependant Variable:	Exporter Component (θ_e)				Buyer Component (θ_b)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
I[Domestic Exp.]	-0.512*	-0.845***	-0.676***	-0.701***	0.064	0.029	-0.002	0.019
	(0.282)	(0.231)	(0.230)	(0.244)	(0.134)	(0.129)	(0.131)	(0.129)
Productive reach Ship.	3	4	5	6	3	4	5	6
Observations	2,082	2,082	2,082	2,082	1,378	1,378	1,378	1,378

Note: This table reports OLS estimates of equation 13 at the exporter-buyer level, i.e., at the relationship level. The dependent variable is the exporter component (Columns 1-4) and the buyer component (Columns 5-8). Columns vary based on the threshold used to define a productive relationship – the number of shipments required to consider a relationship productive – which slightly influences the estimation of θ_e and θ_b in equation 11. Both components are standardized and winsorized at a 5% level. Standard errors, shown in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at a 0.10, 0.05, and 0.01 level, respectively. Observations are weighted so each exporter receives equal weight.

Figure 4: Relationship Ending Controlling for the Buyer

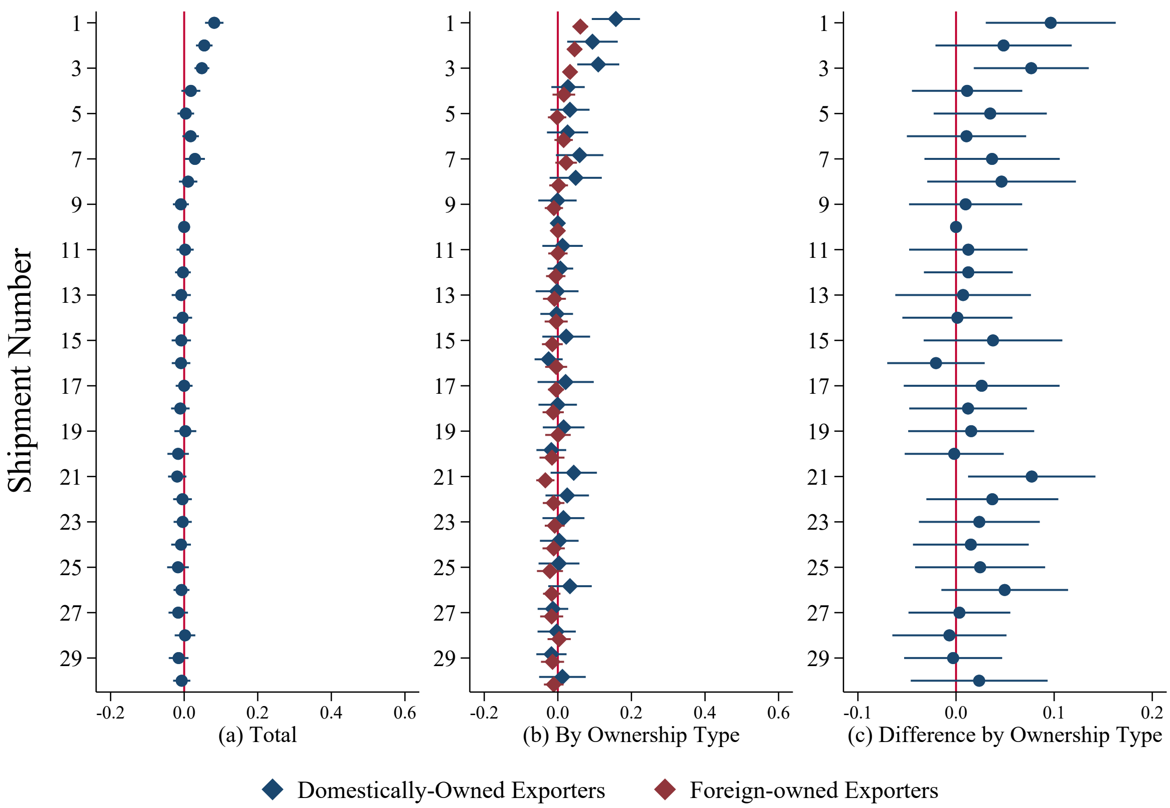

Note: The figure shows the estimated probability and 95% confidence intervals of a buyer-seller relationship ending as a function of the shipment number. Panel (a) plots the estimated coefficients $\hat{\beta}_{1,i}$ from equation 8, which includes buyer fixed effects to account for unobserved heterogeneity in buyer preferences or behavior. Panel (b) displays the coefficient estimates from equation 9, separating the effects for domestically owned exporters ($\hat{\beta}_{1,i}$) and foreign-owned exporters ($\hat{\beta}_{1,i} + \hat{\beta}_{2,i}$), again including buyer fixed effects. Panel (c) illustrates the difference in terminating rates between domestically- and foreign-owned exporters as captured by ($\hat{\beta}_{2,i}$). Standard errors are two-way clustered at the exporter and buyer levels.

Controlling for the buyer substantially reduces early termination of relationships, but it does not eliminate the differences between foreign- and domestically owned exporters (Figure 4). The cumulative probability of a relationship ending within the first three shipments declines by 66% overall – 72% for for foreign and 52% for domestically-owned exporters. Despite this improvement, domestically-owned exporters still face, on average, 12 pp higher probability of relationship termination within the first three shipments compared to the tenth shipment. For foreign-owned exporters, this probability is only 5 pp higher. As a result, a notable disparity persists: the average termination probability for domestically owned exporters remains approximately 7 pp higher than that of foreign-owned exporters during the initial three shipments.

5.3 Exporter Clarity, Relationship Termination and Initiation

Ending Relationships. This section tests our core theoretical prediction: exporters that are more capable of forming productive relationships are also more likely to terminate them in response to positive shocks to their outside option (empirical implication 3). To evaluate this, we estimate

$$Y_{e,t} = \beta_0 + \beta_1 \hat{\theta}_e + \beta_2 Price\ Spread_{t+1} + \beta_3 Price\ Spread_{t+1} \times \hat{\theta}_e + \Gamma X_{e,t} + \epsilon_{e,t}, \tag{14}$$

where the dependent variable, $Y_{e,t}$ is either the number of relationships ended by exporter e in week t, or a indicator variable equal to one if the exporter ended at least one relationship in that week. The key explanatory variable $\hat{\theta}_e$ is the exporter fixed effect from the AKM model, capturing the exporter's ability to form productive relationships. The variable $Price\ Spread_{t+1}$ measures the difference between the average auction price and the average price in direct shipments for the following week. We use the forward-looking spread at $t+1$ since a relationship ending at time t implies that the exporter chose not to ship in $t+1$. As before, we restrict the analysis to productive relationships to isolate the effect of shocks to the exporter's outside option from clarity problems.

Exporters with higher values of θ_e are more likely to terminate productive relationships in response to positive shock to their outside option. For exporters with an average seller clarity component, a one standard deviation increase in the price spread between auctions and direct relationships increases the probability of ending at least one relationship by 0.011 pp (a 21% increase) and the number of relationships that end that week by 0.012 pp (a 20% increase) (Table 3). For exporters with a one standard deviation higher exporter component, $\hat{\theta}_e$, the effect is roughly twice as large: the probability of ending at least one relationship increases by an additional 0.012 pp (23%), and the number of relationships that end rises by an additional 0.013 pp (22%) in the linear specification and by 25% in the Poisson specification.

These findings are robust across a range of alternative specifications, incorporating controls, and definitions of relationship termination. Appendix Table C.9 shows that the results hold when controlling for volume and value of flowers sold by the exporter in a given week. Because extreme prices may influence the price spread, we conduct robustness checks in which we winsorize or trim prices before computing the spread. Appendix Table C.10 confirms that the main results are unchanged when winsorizing prices at the 5% level in both tails.

We test the sensitivity of our results to alternative definitions of relationship termination. Appendix Table C.11 presents consistent estimates based on defining the end of a relationship as no transactions for at least nine months. In this specification, the sample ends three months

Table 3: Exporter Component and Ending Productive Relationships

Dependent Variable:	Number of Relationships Ending ≥ 1 (Dummy)		Number of Relationships Ending (Count)			
	(1)	(2)	(3)	(4)	(5)	(6)
Price Spread (Std)	0.009**	0.011***	0.011**	0.012***	0.028	0.024
	(0.004)	(0.003)	(0.005)	(0.004)	(0.069)	(0.067)
Price Spread (Std) x	0.009**	0.012***	0.010**	0.013***	0.210***	0.221**
Exporter Comp. (θ_e)	(0.003)	(0.003)	(0.004)	(0.004)	(0.072)	(0.086)
Mean Dep. Var	0.052	0.052	0.059	0.059	0.059	0.060
Observations	15,585	15,585	15,585	15,585	15,585	15,329
Exporter FE		Y		Y		Y
Estimation	OLS	OLS	OLS	OLS	PPML	PPML

Note: The table reports OLS and Poisson Pseudo-Maximum Likelihood (PPML) estimates of β_2 and β_3 from equation 14, at the export-week level. In Columns 1-2, the dependent variable is a binary indicator equal to one if the exporter ended at least one relationship during the week. In Columns 3-4, the dependent variable is the number of relationships that ended that week. Columns 5-6 present PPML estimates to account for the count nature of the outcome. A relationship is considered concluded if no shipments are observed between an exporter and a buyer for a period of at least six months. *Price Spread* is measured as the standardized difference between the average auction price and the average price in direct shipments for the following week. The exporter component, $\hat{\theta}_e$, is estimated using equation 11, standardized, and winsorized at a 5% level. The sample is restricted to productive relationships, defined as those that survived beyond the third shipment. All specifications include a control for the number of active, productive relationships the exporter has that week. Standard errors in parentheses are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

earlier to ensure sufficient time for a relationship to meet the termination criteria, i.e., so that included relationships have a chance to end within our sample period. Results are also robust to restricting the analysis to transactions from 2012 onward – when the average price in direct shipments was significantly higher. Appendix Table C.12 shows that although the effect size is larger in this subsample, the restriction reduces our sample size by approximately 30%, also reducing precision. Finally, we verify that the results hold when aggregating the price spread and relationship termination at the monthly level. Appendix Table C.13 reports consistent effects, with larger coefficients due to the aggregation of the dependent variable.

Attempting Relationships. We now test whether exporters that are more adept at forming productive relationships are also more likely to attempt new relationships (empirical implication 2), using the following estimating equation

$$Y_{e,t} = \beta_0 + \beta_1 \hat{\theta}_e + \Gamma X_{e,t} + \zeta_t + \epsilon_{e,t} \tag{15}$$

where the outcome variable, $Y_{e,t}$ is either the number of relationships attempted by exporter e in week t, or a indicator variable equal to one if the exporter attempted at least one relationship during that week. An *attempt* is defined as the first direct transaction between an exporter and a buyer. The exporter component, $\hat{\theta}_e$, measures the ease with which exporters form productive relationships and it is estimated using the AKM framework. To account for time-varying shocks common to all exporters, we include week-by-year fixed effects, ζ_t. The set of controls, $X_{e,t}$, varies across specifications and includes the log quantity and log value of the exporters' transactions in that week.

Table 4: Exporter Component of Task Clarity and Attempting New Relationships

Dependent Variable:	Number of Attempts ≥ 1 (Dummy)			Number of Attempts (Count)					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Exporter Comp. (θ_e)	0.035**	0.036**	0.035**	0.047**	0.048**	0.047**	0.370**	0.377***	0.367**
	(0.016)	(0.016)	(0.016)	(0.023)	(0.023)	(0.023)	(0.146)	(0.145)	(0.148)
Mean Dep. Var	0.105	0.105	0.105	0.127	0.127	0.127	0.142	0.142	0.142
Observations	16,408	16,408	16,408	16,408	16,408	16,408	14,614	14,614	14,614
Controls:									
Quantity (Ln, no. of stems)		Y	Y		Y	Y		Y	Y
Value (Ln, USD)			Y			Y			Y
Week x Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Estimation	OLS	OLS	OLS	OLS	OLS	OLS	PPML	PPML	PPML

Note: The table reports estimates from equation 15, using exporter-week level data. An attempt is defined as the first shipment between an exporter and a buyer. In Columns 1-3, the dependent variable is a binary variable for whether the exporter made at least one attempt that week. In Columns 4-9, the dependent variable is the number of weekly attempts. Columns 7-9 use a Poisson Pseudo-Maximum Likelihood (PPML) estimator. The exporter component, $\hat{\theta_e}$, is estimated using equation 11, then standardized, and winsorized at a 5% level. Controls vary by specifications and include the quantity and value of the exporters' transactions in the week (in logs). Standard errors in parentheses are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Exporters with higher $\hat{\theta_e}$ values – those more likely to reach productive relationships – are indeed more likely to attempt new relationships. A one standard deviation increase in $\hat{\theta_e}$ is associated with 0.05 (37%) more weekly attempts and a 0.04 (33%) increase in the probability of making at least one attempt (Table 4). These results are consistent across all nine specifications. We observe a similar pattern among domestically owned firms, which, on average, struggle more to reach productive relationships. Appendix Table C.14 shows that domestically owned exporters make fewer attempts than their foreign-owned counterparts.

6 Concluding Remarks

We propose and test a model of relationships between a principal and an agent who face a clarity problem. Loosely speaking, clarity is the probability that the agent identifies a productive action for the principal. We show empirically that this problem is not only real but also severe for certain firms. We also document a reversal in termination rates: higher clarity exporters are less likely to end relationships early but more likely to end productive relationships later in response to positive shocks to their outside option.

Our results speak to contemporary debates on industrial policy (e.g., Juhász et al., 2023). Traditional industrial policies – such as land allocation, long-term credit, and logistics subsidies – are "hard," supply-side interventions that require substantial fiscal capacity (Rodrik, 2004). These measures are often effective for kick-starting production and may suffice for non-differentiated products (e.g. steel nuggets or cotton bales), where clarity is less of an issue due to established pricing benchmarks and standardized quality.

However, for differentiated products – where the specifications of the product vary by buyer – clarity problems arise. Since exporting differentiated products directly to buyers is associated with higher seller profitability (Cajal-Grossi et al., 2023) and higher GDP per capita (Rauch, 1999), resolving this problem is critical for moving up the ladder of development. Our model suggests that "soft" interventions, such as export promotion agencies screening buyers or subsidizing trade fairs and origin trips, can help firms overcome clarity challenges. These targeted, lower-cost interventions may be essential for supporting differentiated product exporters.

References

Abowd, John M, Francis Kramarz, and David N Margolis. 1999. "High wage workers and high wage firms." *Econometrica*, 67(2): 251–333.

Adam, Klaus. 2001. "Learning while searching for the best alternative." *Journal of Economic Theory*, 101(1): 252–280.

Alfaro-Ureña, Alonso, Isabela Manelici, and Jose P Vasquez. 2022. "The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages." *Quarterly Journal of Economics*, 137(3): 1495–1552.

Angelucci, Charles, and Roi Orzach. 2025. "Multi-Project Collaborations." MIT Sloan School of Management Working Paper.

Atkin, David, Amit K. Khandelwal, and Adam Osman. 2017. "Exporting and Firm Performance: Evidence from a Randomized Experiment." *Quarterly Journal of Economics*, 132(2): 551–615.

Atkin, David, Arnaud Costinot, and Masao Fukui. 2024. "Globalization and the Ladder of Development: Pushed to the Top or Held at the Bottom?" National Bureau of Economic Research NBER Working Paper 29500.

Auster, Sarah, and Yeon-Koo Che. 2025. "A Theory of Choice Overload." *Working Paper*.

Baker, George, Robert Gibbons, and Kevin J. Murphy. 1994. "Subjective Performance Measures in Optimal Incentive Contracts." *The Quarterly Journal of Economics*, 109(4): 1125–1156.

Baker, George, Robert Gibbons, and Kevin J. Murphy. 2002. "Relational Contracts and the Theory of the Firm." *Quarterly Journal of Economics*, 117(1): 39–84.

Banerjee, Abhijit, and Esther Duflo. 2000. "Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry." *Quarterly Journal of Economics*, 115(3): 989–1017.

Barron, Daniel, Robert Gibbons, Ricard Gil, and Kevin J. Murphy. 2020. "Relational Adaptation Under Reel Authority." *Management Science*, 66(5): 1868–1889.

Bastos, Paulo, Joana Silva, and Eric Verhoogen. 2018. "Export Destinations and Input Prices: Evidence from Portugal." *American Economic Review*, 108(2): 353–392.

Board, Simon. 2011. "Relational contracts and the value of loyalty." *American Economic Review*, 101(7): 3349–3367.

Board, Simon, and Moritz Meyer-ter Vehn. 2015. “Relational Contracts in Competitive Labor Markets.” *Review of Economic Studies*, 82(2): 490—-534.

Bonatti, Alessandro, and Johannes Hörner. 2011. “Collaborating.” *American Economic Review*, 101(2): 632–663.

Bonatti, Alessandro, and Johannes Hörner. 2017. “Career concerns with exponential learning.” *Theoretical Economics*, 12(1): 425–475.

Brugués, Felipe. 2024. “Take the Goods and Run: Contracting Frictions and Market Power in Supply Chains.” *American Economic Review.* Forthcoming.

Cajal-Grossi, Julia, Rocco Macchiavello, and Guillermo Noguera. 2023. “Buyers’ Sourcing Strategies and Suppliers’ Markups in Bangladeshi Garments.” *The Quarterly Journal of Economics*, 138(4): 2391–2450.

Card, David, Jörg Heining, and Patrick Kline. 2013. “Workplace Heterogeneity and the Rise of West German Wage Inequality.” *Quarterly Journal of Economics*, 128(3): 967–1015.

Chade, Hector, Jan Eeckhout, and Lones Smith. 2017. “Sorting through search and matching models in economics.” *Journal of Economic Literature*, 55(2): 493–544.

Chassang, Sylvain. 2010. “Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts.” *American Economic Review*, 100(1): 448–65.

Chawla, Shuchi, Evangelia Gergatsouli, Yifeng Teng, Christos Tzamos, and Ruimin Zhang. 2020. “Pandora’s box with correlations: Learning and approximation.” 1214–1225, IEEE.

Deb, Joyee, Jin Li, and Arijit Mukherjee. 2016. “Relational contracts with subjective peer evaluations.” *The RAND Journal of Economics*, 47(1): 3–28.

Demir, Banu, Ana Cecilia Fieler, Daniel Yi Xu, and Kelly Kaili Yang. 2024. “O-Ring Production Networks.” *Journal of Political Economy*, 132(1): 200–247.

Fahn, Matthias, W. Bentley MacLeod, and Gerd Mühlheusser. 2023. “Symposium on the Economics of Relational Contracts: Past and Future Developments.” *Journal of Institutional and Theoretical Economics*, 179(3–4): 413–440.

Fong, Yuk-Fai, and Jin Li. 2016. “Information revelation in relational contracts.” *The Review of economic studies*, 84(1): 277–299.

Fuchs, William. 2007. “Contracting with repeated moral hazard and private evaluations.” *American Economic Review*, 97(4): 1432–1448.

Gergatsouli, Evangelia, and Christos Tzamos. 2023. “Weitzman’s rule for pandora’s box with correlations.” *Advances in Neural Information Processing Systems*, 36: 12644–12664.

Ghosh, Parikshit, and Debraj Ray. 1996. “Cooperation in Community Interaction Without Information Flows.” *The Review of Economic Studies*, 63(3): 491–519.

Gibbons, Robert, and Rebecca Henderson. 2012. “Relational Contracts and Organizational Capabilities.” *Organization Science*, 23(5): 1350–1364.

Gil, Ricard, Myongjin Kim, and Giorgio Zanarone. 2022. “Relationships Under Stress: Relational Outsourcing in the U.S. Airline Industry After the 2008 Financial Crisis.” *Management Science*, 68(2): 1256–1277.

Halac, Marina. 2012. “Relational Contracts and the Value of Relationships.” *American Economic Review*, 102(2): 750–79.

Halac, Marina. 2014. “Relationship building: conflict and project choice over time.” *The Journal of Law, Economics, & Organization*, 30(4): 683–708.

Hansman, Christopher, Jonas Hjort, Gianmarco Leon, and Matthieu Teachout. 2020. “Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters.” *Journal of Political Economy*, 128(9): 3570–3625.

Harris, Adam, and Thi Mai Anh Nguyen. 2025a. “Long-Term Relationships and the Spot Market: Evidence from US Trucking.” New York University.

Harris, Adam, and Thi Mai Anh Nguyen. 2025b. “Long-Term Relationships in the US Truckload Freight Industry.” *American Economic Journal: Microeconomics*, 17(1): 308–53.

Janssen, Maarten, Thomas Jungbauer, Marcel Preuss, and Cole Randall Williams. 2025. “Search Platforms: Big Data and Sponsored Positions.” Working Paper.

Jordan v. City of New London. 2000. “Jordan v. City of New London.” https://www.aele.org/apa/jordan-newlondon.html, Accessed: 2025-05-12.

Jovanovic, Boyan. 1979. “Job matching and the theory of turnover.” *Journal of Political Economy*, 87(5): 972–990.

Jovanovic, Boyan. 1984. “Matching, turnover, and unemployment.” *Journal of political Economy*, 92(1): 108–122.

Juhász, Réka, Nathan J Lane, and Dani Rodrik. 2023. “The New Economics of Industrial Policy.” , (31538).

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. 2020. “Leave-out estimation of variance components.” *Econometrica*, 88(5): 1859–1898.

Kostadinov, Rumen, and Aditya Kuvalekar. 2022. “Learning in relational contracts.” *American Economic Journal: Microeconomics*, 14(1): 284–329.

Kranton, Rachel E. 1996. “The formation of cooperative relationships.” *The Journal of Law, Economics, and Organization*, 12(1): 214–233.

Lazear, Edward P. 2009. “Firm-specific human capital: A skill-weights approach.” *Journal of political economy*, 117(5): 914–940.

Levin, Jonathan. 2003. “Relational Incentive Contracts.” *American Economic Review*, 93(3): 835–857.

Macchiavello, Rocco. 2022. “Relational Contracts and Development.” *Annual Review of Economics*, 14: 337–362.

Macchiavello, Rocco, and Ameet Morjaria. 2015. "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports." *American Economic Review*, 105(9): 2911–45.

Macchiavello, Rocco, and Ameet Morjaria. 2021. "Competition and Relational Contracts in the Rwanda Coffee Chain." *Quarterly Journal of Economics*, 136(2): 1089–1143.

Macchiavello, Rocco, and Ameet Morjaria. 2023. "Relational Contracts: Recent Empirical Advancements and Open Questions." *Journal of Institutional and Theoretical Economics*, 179(3): 673 – 700.

MacLeod, W. Bentley. 2007. "Reputations, Relationships, and Contract Enforcement." *Journal of Economic Literature*, 45(3): 595–628.

Macleod, W. Bentley, and James Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica*, 57(2): 447–80.

Malcomson, James M. 2013. "Relational Incentive Contracts." In The Handbook of Organizational Economics. , ed. Robert Gibbons and John Roberts, 1014–1065. Princeton University Press.

Mano, Yukichi, Takashi Yamano, Aya Suzuki, and Tomoya Matsumoto. 2011. "Local and Personal Networks in Employment and the Development of Labor Markets: Evidence from the Cut Flower Industry in Ethiopia." *World Development*, 39(10): 1760–1770.

McAdams, David. 2011. "Performance and turnover in a stochastic partnership." *American Economic Journal: Microeconomics*, 3(4): 107–142.

McMillan, John, and Christopher Woodruff. 1999. "Interfirm Relationships and Informal Credit in Vietnam." *Quarterly Journal of Economics*, 114(4): 1285–1320.

Mortensen, Dale T, and Christopher A Pissarides. 1999. "New developments in models of search in the labor market." *Handbook of labor economics*, 3: 2567–2627.

Moscarini, Giuseppe. 2005. "Job matching and the wage distribution." *Econometrica*, 73(2): 481–516.

Nocke, Volker, and Patrick Rey. 2024. "Consumer search, steering, and choice overload." *Journal of Political Economy*, 132(5): 1684–1739.

Oqubay, Arkebe. 2015. Made in Africa: Industrial Policy in Ethiopia. Oxford University Press.

Quah, John K-H, and Bruno Strulovici. 2013. "Discounting, values, and decisions." *Journal of Political Economy*, 121(5): 896–939.

Rauch, James E. 1999. "Networks versus markets in international trade." *Journal of International Economics*, 48(1): 7–35.

Rodrik, Dani. 2004. "Industrial Policy for the Twenty-First Century." *Working Paper, Harvard.*

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. 2019. "Firming up inequality." *Quarterly Journal of Economics*, 134(1): 1–50.

Verhoogen, Eric. 2008. "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector." *Quarterly Journal of Economics*, 123(2): 489–530.

Volpe Martincus, Christian, and Jeronimo Carballo. 2008. "Is Export Promotion Effective in Developing Countries? Firm-level Evidence on the Intensive and the Extensive Margins of Exports." *Journal of International Economics*, 76(1): 89–106.

Watson, Joel. 1999. "Starting small and renegotiation." *Journal of Economic Theory*, 85(1): 52–90.

Weitzman, Martin L. 1979. "Optimal search for the best alternative." *Econometrica*, 47(3): 641.

Wright, Randall, Philipp Kircher, Benoît Julien, and Veronica Guerrieri. 2021. "Directed search and competitive search equilibrium: A guided tour." *Journal of Economic Literature*, 59(1): 90–148.

Appendix

A Optimal search for a productive action

This Appendix provides a search-based microfoundation for Assumption 1. Suppose $|\mathcal{A}_i| = N$. A state in the search problem is a vector which describes whether each of the N actions is productive or not. Let $\mathcal{P} \subset \mathcal{A}_i$ denote the set of productive actions in a state. Assume that either action a_k is productive with probability π_k independent of other actions, i.e., the joint distribution is the product of the marginals described by π_k, or none of the actions are productive, i.e., the joint distribution assigns probability 1 to state where $a_k \notin \mathcal{P}$, for all k. The joint distribution in our search problem is an *affine* combination which puts weight $q > 0$ on the independent distribution and weight $1 - q$ on none of the actions being productive.

If $q = 1$ we have an independent distribution with marginals π_k. If $q < 1$ there is symmetric positive correlation, so that learning one action is unproductive increases the likelihood that none are. This can arise if the principal might be "impossible to please". If $q > 1$ we have negative correlation and the distribution is well-defined as long as $q \leq 1/\left(1 - \mathbb{P}\left[a_k \notin \mathcal{P}, \forall k\right]\right)$. At the upper-bound for q the agent is certain that at least one action is productive.

When $q = 1$ actions are independent and, following the logic of Weitzman (1979), the agent optimally searches in descending order of π_k, since all boxes carry the same "prize". We can thus reorder actions so that $\pi_0 \geq \pi_1 \geq ... \geq \pi_{N-1}$. Even when $q \neq 1$, the agent still optimally searches in descending order of π_k, since the correlation is symmetric.

The conditional probability of finding a productive action on the agent's k^{th} attempt, given that no productive action has been found is

$$\lambda_k \quad = \quad \mathbb{P}\left[a_k \in \mathcal{P} \mid \forall j < k, a_j \notin \mathcal{P}\right] = \pi_k \frac{q \prod_{j=0}^{k-1} (1 - \pi_j)}{q \prod_{j=0}^{k-1} (1 - \pi_j) + (1 - q)}.$$

When $q = 1$, $\lambda_k = \pi_k$ and $(\lambda_k)_{k=0}^{N-1}$ is decreasing. When $q \in (0, 1)$, this sequence is also decreasing. To see this, observe that π_k is decreasing in k and that the second factor above is also decreasing in k when $q < 1$.

With negative correlation, $q > 1$, and we may get single-peaked sequences of $(\lambda_k)_{k=0}^{N-1}$ since the second factor is increasing in k. To give a specific example, consider $N = 4$ with $\pi_0 = 0.55$, $\pi_1 = 0.5$, $\pi_2 = 0.3$, $\pi_3 = 0.05$. While the optimal search strategy is in descending order of π_k, λ_k need not be decreasing; it can be single-peaked. Table A.1 shows the sequence $(\lambda_k)_{k=0}^{3}$ for several $q \geq 1$.

Table A.1: $(\lambda_k)_{k=0}^{3}$ with negative correlation

q	λ_0	λ_1	λ_2	λ_3
1	<u>0.55</u>	0.50	0.30	0.05
1.05	<u>0.5775</u>	0.5592	0.3805	0.0717
1.1	0.6050	<u>0.6266</u>	0.5034	0.1183
1.15	0.6325	0.7041	<u>0.7138</u>	0.2910
1.175	0.6463	0.7473	0.8874	<u>0.9196</u>

Note: This table shows $(\lambda_k)_{k=0}^{3}$ for different choices of $q \geq 1$. All of them are single-peaked, with the peak underlined.

The value $q = 1.175$ is close to the upper bound and $\mathbb{P}[a_k \notin \mathcal{P}, \forall k]$ is close to 0. Thus, even though a_3 is the least likely to be productive, because the agent is very sure that at least one action must be productive, λ_3 is close to 1.

The setting and question we ask is novel in the search literature. While we add complexity by wanting to tackle correlation, our problem is simpler because all prizes are equal and so the stopping rule is trivial. Furthermore, we are not interested in finding an index policy (which has been the focus of this literature) but are instead interested in characterizing the conditional probability of finding a prize. Although little is known about the Pandora's box problem with correlation, numerical analysis suggests that these problems can often result in an optimal search sequence where $(\lambda_k)_{k=0}^{N-1}$ is single-peaked, even when correlation is not symmetric. While one can easily construct examples where this does not hold, a characterization of the primitives of the search problem that lead to a single-peaked $(\lambda_k)_{k=0}^{N-1}$ is beyond the scope of this paper.[36]

Assumption 1 is more general than the microfoundation given above. For example, a model where all actions are productive with the same, but unknown, probability, and the agent learns about this probability, also satisfies Assumption 1. Note that Assumption 1 extends the sequence of $(\lambda_t)_{t=0}^{\infty}$ to be infinite. If $|\mathcal{A}_i| = N < \infty$, we set $\lambda_k = 0$ for all $k \geq N$. We also assume that $(\lambda_t)_{t=0}^{\infty}$ converges to zero, i.e., after a large number of failures to find a productive action, the agent's belief that he will find one on the next try is close to zero. This rules out the trivial case where the agent never finds a productive action, but keeps trying forever.

B Omitted Proofs

Lemma 1. *Fix W_0. There exists a unique V which satisfies equation (2). This V is increasing in W_0. Furthermore the optimal cutoff shock s^* is decreasing in W_0.*

[36] For example, set $\pi_1 = 0.31$ and take q to be sufficiently high in the above example. We can derive conditions on $(\pi_k)_{k=0}^{N-1}$ to guarantee $(\lambda_k)_{k=0}^{N-1}$ is single-peaked for all q. With three actions it suffices that $(1-\pi_0)\pi_1 > (1-\pi_1)\pi_2$. Cutoffs for q can then be derived that replicate the pattern in Table A.1.

Proof. Equation (2) can be written as $V = \int_\ell^\infty \max\{b + \delta V, s + \delta W_0 - c\} \, \mathrm{d}F(s)$. Define the operator $T : \mathbb{R} \to \mathbb{R}$ by

$$T(V) = \int_\ell^\infty \max\{b + \delta V, s + \delta W_0 - c\} \, \mathrm{d}F(s).$$

We now verify Blackwell's sufficient conditions for a contraction mapping. Clearly T is monotone, since for any $V' \geq V$, $\max\{b + \delta V', s + \delta W_0 - c\} \geq \max\{b + \delta V, s + \delta W_0 - c\}$ and integration with respect to F preserves this monotonicity. To see that T satisfies discounting, take any $\alpha \geq 0$ and note that $\max\{b + \delta(V + \alpha), s + \delta W_0 - c\} \leq \max\{b + \delta V, s + \delta W_0 - c\} + \delta\alpha$. Integrating with respect to F, yields

$$T(V + \alpha) = \int_\ell^\infty \max\{b + \delta(V + \alpha), s + \delta W_0 - c\} \, \mathrm{d}F(s) \leq T(V) + \delta\alpha.$$

Thus T satisfies Blackwell's sufficient conditions and is a contraction mapping with modulus δ.

For a fixed W_0 and corresponding fixed point V, let $s^* = \delta (V - W_0) + b + c$ be the shock that makes the agent indifferent between the outside option and continuing in the relationship. Equation (2) implicitly defines V as follows

$$V(W_0) = \int_\ell^{s^*} (b + \delta V(W_0)) \ \mathrm{d}F(s) + \int_{s^*}^\infty (s + \delta W_0 - c) \ \mathrm{d}F(s).$$

Differentiating both sides with respect to W_0 we have that

$$\frac{\mathrm{d}V}{\mathrm{d}W_0} = \int_\ell^{s^*} \delta \frac{\mathrm{d}V}{\mathrm{d}W_0} \ \mathrm{d}F(s) + \int_{s^*}^\infty \delta \ \mathrm{d}F(s) = \delta F(s^*) \frac{\mathrm{d}V}{\mathrm{d}W_0} + \delta (1 - F(s^*)).$$

Note that the terms involving the limits of the integrals cancel since the functions coincide at $s = s^*$ by definition of s^*. Solving for $\frac{\mathrm{d}V}{\mathrm{d}W_0}$ we find

$$\frac{\mathrm{d}V}{\mathrm{d}W_0} = \delta \frac{1 - F(s^*)}{1 - \delta F(s^*)} \in [0, \delta].$$

We have that $\frac{\mathrm{d}V}{\mathrm{d}W_0} > 0$ if and only if $F(s^*) < 1$. This is intuitive since if $F(s^*) = 1$, the agent almost never terminates a productive relationship and $V = \frac{b}{1-\delta}$, which is independent of W_0. Observe further that $F(s^*) > 0$ and $\frac{\mathrm{d}V}{\mathrm{d}W_0} < \delta$, if the agent is to stay in a productive relationship with positive probability; this is necessary for (RC) to hold for any $c > 0$.

Finally, since $s^* = \delta (V - W_0) + b + c$, we have that $\frac{\mathrm{d}s^*}{\mathrm{d}W_0} = \delta \left(\frac{\mathrm{d}V}{\mathrm{d}W_0} - 1 \right) \leq \delta (\delta - 1) < 0$. □

B.1 Proof of Theorem 1

We make a useful observation. After some number of failed attempts to find a productive action, the agent terminates the relationship since $\lambda_n (b + \delta V) + (1 - \lambda_n)(\delta W_0 - c) < \ell + \delta W_0 - c$, because λ_n converges to 0 as $n \to \infty$. Period n is thus an upper bound on how long an unproductive relationship can last. This upper bound will be used to iterate value functions backward. If $\mathcal{A}_i$ is finite, we can take $n = |\mathcal{A}_i|$. As a convention, we let $W_t = 0$ for all $t > n$, but these periods will not be relevant for us.

Theorem. *A unique solution to the agent's problem exists where the agent attempts to form* *relationships as long as (RC) holds.*

Let $T_0 : \mathbb{R} \to \mathbb{R}$ be this map that takes a guess for W_0 and maps it to another W_0. The closed form expression is not simple in general (although one can get it for special cases, e.g., if F puts probability 1 on a single point). To show that T_0 is a contraction mapping take two guesses for W_0, namely $W_0^+ > W_0^-$ and we will show that $T_0\left(W_0^+\right) - T_0\left(W_0^-\right) < W_0^+ - W_0^-$. We will denote by $W_k\left(s_t|W_0\right)$, the value of $W_k\left(s_t\right)$ when starting with a particular W_0 and iterating backwards. First observe that by equation (2), $0 \le V\left(W_0^+\right) - V\left(W_0^-\right)$ and that

$$V\left(W_0^+\right) - V\left(W_0^-\right) \le \delta\left(W_0^+ - W_0^-\right),$$

since $V'\left(W_0\right) \le \delta$ by Lemma 1.

Recall that $W_t\left(s_t\right) = \max\left\{s_t + \delta W_0 - c, \lambda_t\left(b + \delta V\right) + \left(1-\lambda_t\right)\max\left\{\delta W_{t+1}, \delta W_0 - c\right\}\right\}$. Since a relationship ends before period n, we have that $W_n\left(W_0^+\right) = \upsilon + \delta W_0^+ - c$ and $W_n\left(W_0^-\right) = \upsilon + \delta W_0^- - c$. As such

$$W_n\left(W_0^+\right) - W_n\left(W_0^-\right) \le \delta\left(W_0^+ - W_0^-\right). \tag{16}$$

Lemma 2. *If* $W_{t+1}\left(W_0^+\right) - W_{t+1}\left(W_0^-\right) \le \delta\left(W_0^+ - W_0^-\right)$ *and* $\max\left\{\delta W_{t+1}\left(W_0^-\right), \delta W_0^- - c\right\} = \delta W_0^- - c$, *then* $\max\left\{\delta W_{t+1}\left(W_0^+\right), \delta W_0^+ - c\right\} = \delta W_0^+ - c$.

Proof. Assume by way of contradiction that $\delta W_{t+1}\left(W_0^-\right) < \delta W_0^- - c$ and that $\delta W_{t+1}\left(W_0^+\right) > \delta W_0^+ - c$. Subtracting the first equation from the second results in $\delta W_{t+1}\left(W_0^+\right) - \delta W_{t+1}\left(W_0^-\right) > \delta\left(W_0^+ - W_0^-\right)$, but this contradicts that $W_{t+1}\left(W_0^+\right) - W_{t+1}\left(W_0^-\right) \le \delta\left(W_0^+ - W_0^-\right)$. □

Lemma 3. *Suppose that* $W_{t+1}\left(W_0^+\right) - W_{t+1}\left(W_0^-\right) \le \delta\left(W_0^+ - W_0^-\right)$. *Then* $W_t\left(s_t|W_0^-\right) = s_t + \delta W_0^- - c$, *implies* $W_t\left(s_t|W_0^+\right) = s_t + \delta W_0^+ - c$.

Proof. We consider two cases, either $\delta W_{t+1}\left(W_0^+\right) \le \delta W_0^+ - c$ or the opposite.

Case 1: $\max\left\{\delta W_{t+1}\left(W_0^+\right), \delta W_0^+ - c\right\} = \delta W_0^+ - c$.

Since $W_t\left(s_t|W_0^-\right) = s_t + \delta W_0^- - c$, we have

$$\begin{aligned} s_t + \delta W_0^- - c &\ge \lambda_t\left(b + \delta V\left(W_0^-\right)\right) + \left(1-\lambda_t\right)\max\left\{\delta W_{t+1}\left(W_0^-\right), \delta W_0^- - c\right\} \\ &\ge \lambda_t\left(b + \delta V\left(W_0^-\right)\right) + \left(1-\lambda_t\right)\left(\delta W_0^- - c\right). \end{aligned}$$

This implies

$$0 \le s_t - \lambda_t\left(b + \delta V\left(W_0^-\right)\right) - \lambda_t c + \lambda_t \delta W_0^- \le s_t - \lambda_t\left(b + \delta V\left(W_0^+\right)\right) - \lambda_t c + \lambda_t \delta W_0^+, \tag{17}$$

where the second inequality follows since

$$\lambda_t \delta W_0^+ - \lambda_t \delta W_0^- - \lambda_t\left(\delta V\left(W_0^+\right) - \delta V\left(W_0^-\right)\right) \ge \lambda_t \delta\left(W_0^+ - W_0^-\right) - \lambda_t \delta^2\left(W_0^+ - W_0^-\right) \ge 0.$$

Rearranging inequality (17) yields $s_t + \delta W_0^+ - c \ge \lambda_t\left(b + V\left(W_0^+\right)\right) + \left(1-\lambda_t\right)\left(\delta W_0^+ - c\right)$, which is what we wanted to show.

Case 2: $\max\left\{\delta W_{t+1}\left(W_0^+\right), \delta W_0^+ - c\right\} = \delta W_{t+1}\left(W_0^+\right)$.

From Lemma 2, we have that $\max\left\{\delta W_{t+1}\left(W_0^-\right), \delta W_0^- - c\right\} = \delta W_{t+1}\left(W_0^-\right)$. So $s_t + \delta W_0^- - c \geq \lambda_t\left(b + \delta V\left(W_0^-\right)\right) + (1-\lambda_t)\,\delta W_{t+1}\left(W_0^-\right)$, holds only if $\lambda_t\left(b + \delta V\left(W_0^-\right)\right) + (1-\lambda_t)\,\delta W_{t+1}\left(W_0^-\right) - \delta W_0^- \leq s_t - c$. Since

$$\begin{aligned} & \lambda_t\left(\delta V\left(W_0^+\right) - \delta V\left(W_0^-\right)\right) + (1-\lambda_t)\,\delta\left(W_{t+1}\left(W_0^+\right) - W_{t+1}\left(W_0^-\right)\right) \\ \leq\ & \lambda_t\delta^2\left(W_0^+ - W_0^-\right) + (1-\lambda_t)\,\delta^2\left(W_0^+ - W_0^-\right) < \delta\left(W_0^+ - W_0^-\right), \end{aligned}$$

$\lambda_t\left(b + \delta V\left(W_0^+\right)\right) + (1-\lambda_t)\,\delta W_{t+1}\left(W_0^+\right) - \delta W_0^+ < \lambda_t\left(b + V\left(W_0^-\right)\right) + (1-\lambda_t)\,\delta W_{t+1}\left(W_0^-\right) - \delta W_0^-$. Thus $\lambda_t\left(b + \delta V\left(W_0^+\right)\right) + (1-\lambda_t)\,\delta W_{t+1}\left(W_0^+\right) - \delta W_0^+ < s_t - c$. This implies that $s_t + \delta W_0^+ - c > \lambda_t\left(b + \delta V\left(W_0^+\right)\right) + (1-\lambda_t)\,\delta W_{t+1}\left(W_0^+\right)$, which is what we wanted to show. □

Lemma 4. *Suppose that* $W_{t+1}\left(W_0^+\right) - W_{t+1}\left(W_0^-\right) \leq \delta\left(W_0^+ - W_0^-\right)$. *Then* $W_t\left(s_t|W_0^+\right) - W_t\left(s_t|W_0^-\right) \leq \delta\left(W_0^+ - W_0^-\right)$.

Proof. Recall that $W_t(s_t) = \max\{s_t + \delta W_0 - c, \lambda_t(b + \delta V) + (1-\lambda_t)\max\{\delta W_{t+1}, \delta W_0 - c\}\}$. We consider two cases.

Case 1: $W_t\left(s_t|W_0^+\right) = s_t + \delta W_0^+ - c$.

By definition $W_t\left(s_t|W_0^-\right) \geq s_t + \delta W_0^- - c$, and

$$W_t\left(s_t|W_0^+\right) - W_t\left(s_t|W_0^-\right) \leq s_t + \delta W_0^+ - c - s_t - \delta W_0^- + c = \delta\left(W_0^+ - W_0^-\right).$$

Case 2: $W_t\left(s_t|W_0^+\right) = \lambda_t\left(b + \delta V\left(W_0^+\right)\right) + (1-\lambda_t)\max\left\{\delta W_{t+1}\left(W_0^+\right), \delta W_0^+ - c\right\}$.

By Lemma 3, $W_t\left(s_t|W_0^-\right) = \lambda_t\left(b + \delta V\left(W_0^-\right)\right) + (1-\lambda_t)\max\left\{\delta W_{t+1}\left(W_0^-\right), \delta W_0^- - c\right\}$. We will then break it up into three further sub-cases.

Case 2a: $\max\left\{\delta W_{t+1}\left(W_0^-\right), \delta W_0^- - c\right\} = \delta W_0^- - c$.

By Lemma 2, $\max\left\{\delta W_{t+1}\left(W_0^+\right), \delta W_0^+ - c\right\} = \delta W_0^+ - c$. Thus

$$\begin{aligned} W_t\left(s_t|W_0^+\right) - W_t\left(s_t|W_0^-\right) &= \lambda_t\left(b + \delta V\left(W_0^+\right)\right) + (1-\lambda_t)\left(\delta W_0^+ - c\right) \\ &\quad -\lambda_t\left(b + \delta V\left(W_0^-\right)\right) - (1-\lambda_t)\left(\delta W_0^- - c\right) \\ &= \lambda_t\delta\left(V\left(W_0^+\right) - V\left(W_0^-\right)\right) + (1-\lambda_t)\,\delta\left(W_0^+ - W_0^-\right) \\ &< \delta\left(W_0^+ - W_0^-\right). \end{aligned}$$

Case 2b: $\max\left\{\delta W_{t+1}\left(W_0^-\right), \delta W_0^- - c\right\} = \delta W_{t+1}\left(W_0^-\right)$ and $\max\left\{\delta W_{t+1}\left(W_0^+\right), \delta W_0^+ - c\right\} = \delta W_{t+1}\left(W_0^+\right)$. So that

$$\begin{aligned} & W_t\left(s_t|W_0^+\right) - W_t\left(s_t|W_0^-\right) \\ =\ & \lambda_t\left(b + \delta V\left(W_0^+\right)\right) + (1-\lambda_t)\,\delta W_{t+1}\left(W_0^+\right) - \lambda_t\left(b + \delta V\left(W_0^-\right)\right) - (1-\lambda_t)\,\delta W_{t+1}\left(W_0^-\right) \\ \leq\ & \lambda_t\delta^2\left(W_0^+ - W_0^-\right) + (1-\lambda_t)\,\delta\left(W_{t+1}\left(W_0^+\right) - W_{t+1}\left(W_0^-\right)\right) \\ \leq\ & \lambda_t\delta\left(W_0^+ - W_0^-\right) + (1-\lambda_t)\,\delta^2\left(W_0^+ - W_0^-\right) < \delta\left(W_0^+ - W_0^-\right) \end{aligned}$$

Case 2c: $\max\left\{\delta W_{t+1}\left(W_0^-\right), \delta W_0^- - c\right\} = \delta W_{t+1}\left(W_0^-\right)$ and $\max\left\{\delta W_{t+1}\left(W_0^+\right), \delta W_0^+ - c\right\} = \delta W_0^+ - c$. Now

$$\begin{aligned} & W_t\left(s_t|W_0^+\right) - W_t\left(s_t|W_0^-\right) \\ = \ & \lambda_t\left(b + \delta V\left(W_0^+\right)\right) + (1-\lambda_t)\left(\delta W_0^+ - c\right) - \lambda_t\left(b + \delta V\left(W_0^-\right)\right) - (1-\lambda_t)\,\delta W_{t+1}\left(W_0^-\right) \\ \leq \ & \lambda_t \delta^2\left(W_0^+ - W_0^-\right) + (1-\lambda_t)\left(\delta W_0^+ - c - \delta W_{t+1}\left(W_0^-\right)\right) \\ < \ & \lambda_t \delta\left(W_0^+ - W_0^-\right) + (1-\lambda_t)\left(\delta W_0^+ - c - \delta W_0^- + c\right) = \delta\left(W_0^+ - W_0^-\right), \end{aligned}$$

where the second last line follows since $\delta W_{t+1}\left(W_0^-\right) \geq \delta W_0^- - c$. □

Proof of Theorem: We first show that the map $T_0 : \mathbb{R} \to \mathbb{R}$ is a contraction. Take $W_0^+ > W_0^-$. By equation (16), $W_n\left(W_0^+\right) - W_n\left(W_0^-\right) \leq \delta\left(W_0^+ - W_0^-\right)$. We continue by backward induction. Lemma 4 showed that $W_{t+1}\left(W_0^+\right) - W_{t+1}\left(W_0^-\right) \leq \delta\left(W_0^+ - W_0^-\right)$ implies

$$W_t\left(W_0^+\right) - W_t\left(W_0^-\right) = \int_\ell^\infty \left[W_t\left(s|W_0^+\right) - W_t\left(s|W_0^-\right)\right] \, \mathrm{d}F(s) \leq \delta\left(W_0^+ - W_0^-\right).$$

Continuing till $t = 0$, we have $T_0\left(W_0^+\right) - T_0\left(W_0^-\right) = W_0\left(W_0^+\right) - W_0\left(W_0^-\right) \leq \delta\left(W_0^+ - W_0^-\right)$. By the Banach Fixed-Point Theorem there exists a unique W_0^* which satisfies the above. Thus, the agent's payoff is uniquely defined. The only possible non-uniqueness in the strategy is if the agent is indifferent between making an additional attempt with the current principal or matching with a new one. In this case we break ties in favor of the current principal.

The Lemma below derives a sufficient condition for the relationship constraint (RC). As part of the proof it develops a necessary and sufficient condition which is not a function of the primitives, but which will be useful later.

Lemma 5. *Inequality (RC) holds if* $\lambda_0 \geq \frac{(1-\delta)(\delta v + c)}{\delta(1-\delta)(b+\delta V) - \delta^2 v}$.

Proof. We will focus on the case where

$$\begin{aligned} W_0(s) \ &= \ \lambda_0(b + \delta V) + (1-\lambda_0)(\delta W_0 - c) \\ &\leq \ \max\left\{s + \delta W_0 - c, \lambda_0(b + \delta V) + (1-\lambda_0)\max\{\delta W_1, \delta W_0 - c\}\right\}. \end{aligned}$$

This is a lower-bound on $W_0(s)$ at which $W_0 = \frac{\lambda_0(b+\delta V) - c(1-\lambda_0)}{1-\delta(1-\lambda_0)}$. Thus a sufficient condition for inequality (RC) is

$$\delta\left(\frac{\lambda_0(b+\delta V) - c(1-\lambda_0)}{1-\delta(1-\lambda_0)}\right) - c = \frac{\delta\lambda_0(b+\delta V) - c}{1-\delta(1-\lambda_0)} \geq \frac{\delta v}{1-\delta}.$$

Solving for λ_0 we get

$$\lambda_0 \geq \frac{(1-\delta)(\delta v + c)}{\delta(1-\delta)(b+\delta V) - \delta^2 v}.$$

Observing that $V \geq \frac{b}{1-\delta}$ and simplifying gives the sufficient condition we wanted to prove. □

B.2 Proof of Theorem 2

Theorem. *The sequence of continuation values, $(W_t)_{t=0}^n$ is single peaked with a peak $\tau \leq T$.*

We start with the following lemma.

Lemma 6. *For all $t \geq T$, $W_t \geq W_{t+1}$.*

Proof. We first show that if $W_{t+1} \geq W_{t+2}$ and $\lambda_t \geq \lambda_{t+1}$, then $W_t \geq W_{t+1}$. Recall that

$$W_t(s) = \max\{s + \delta W_0 - c, \lambda_t (b + \delta V) + (1 - \lambda_t) \max\{\delta W_{t+1}, \delta W_0 - c\}\}.$$

1) Suppose that $s > s^*_{t+1}$. Then

$$\begin{aligned} W_t(s) - W_{t+1}(s) &= \max\{s + \delta W_0 - c, \lambda_t (b + \delta V) + (1 - \lambda_t) \max\{\delta W_{t+1}, \delta W_0 - c\}\} - (s + \delta W_0 - c) \\ &\geq s + \delta W_0 - c - (s + \delta W_0 - c) = 0. \end{aligned}$$

2) Suppose that $s \leq s^*_{t+1}$. Then

$$\begin{aligned} W_t(s) - W_{t+1}(s) &= \max\{s + \delta W_0 - c, \lambda_t (b + \delta V) + (1 - \lambda_t) \max\{\delta W_{t+1}, \delta W_0 - c\}\} \\ &\quad - (\lambda_{t+1}(b + \delta V) + (1 - \lambda_{t+1}) \max\{\delta W_{t+2}, \delta W_0 - c\}) \\ &\geq (\lambda_t - \lambda_{t+1}) \delta V + (1 - \lambda_t) \max\{\delta W_{t+1}, \delta W_0 - c\} \\ &\quad - (1 - \lambda_{t+1}) \max\{\delta W_{t+2}, \delta W_0 - c\}. \end{aligned}$$

Now, if $\delta W_{t+2} < \delta W_0 - c$ we get (since $\lambda_t \geq \lambda_{t+1}$ and $V \geq W_0$)

$$\begin{aligned} W_t(s) - W_{t+1}(s) &= (\lambda_t - \lambda_{t+1}) \delta V + (1 - \lambda_t) \max\{\delta W_{t+1}, \delta W_0 - c\} - (1 - \lambda_{t+1})(\delta W_0 - c) \\ &\geq (\lambda_t - \lambda_{t+1}) \delta V + (1 - \lambda_t)(\delta W_0 - c) - (1 - \lambda_{t+1})(\delta W_0 - c) \\ &= (\lambda_t - \lambda_{t+1})(\delta V - \delta W_0 + c) \geq 0. \end{aligned}$$

But if $\delta W_{t+2} \geq \delta W_0 - c$ we have

$$\begin{aligned} W_t(s) - W_{t+1}(s) &= (\lambda_t - \lambda_{t+1}) \delta V + (1 - \lambda_t) \max\{\delta W_{t+1}, \delta W_0 - c\} - (1 - \lambda_{t+1}) \delta W_{t+2} \\ &\geq (\lambda_t - \lambda_{t+1}) \delta V + (1 - \lambda_t) \delta W_{t+1} - (1 - \lambda_{t+1}) \delta W_{t+2} \\ &= (\lambda_t - \lambda_{t+1}) \delta V + (1 - \lambda_t)(\delta W_{t+1} - \delta W_{t+2}) - (1 - \lambda_{t+1}) \delta W_{t+2} + (1 - \lambda_t) \delta W_{t+2} \\ &= (\lambda_t - \lambda_{t+1})(\delta V - \delta W_{t+2}) + (1 - \lambda_t) \delta (W_{t+1} - W_{t+2}) \geq 0, \end{aligned}$$

where the last inequality follows from $\lambda_t \geq \lambda_{t+1}$, $V \geq W_{t+2}$ and $W_{t+1} \geq W_{t+2}$.

Integrating over s we get $W_t - W_{t+1} = \int (W_t(s) - W_{t+1}(s)) \, dF(s) \geq 0$, which proves that if $W_{t+1} \geq W_{t+2}$ and $\lambda_t \geq \lambda_{t+1}$, then $W_t \geq W_{t+1}$.

To complete the proof of the lemma, recall that $W_n = v + \delta W_0 - c$. Hence $W_{n-1} \geq W_n$ trivially. Furthermore, $n > T$ if inequality (RC) holds and hence for $t \geq T$, $\lambda_t \geq \lambda_{t+1}$. Working backwards from n we have that for all $t \geq T$, $W_t \geq W_{t+1}$. □

To prove the proposition, let $\tau = \max\{t \leq T : W_t \geq W_{t'} \text{ for all } t'\}$. Note that this is well defined since Lemma 6 implies that $W_T \geq W_{t'}$ for all $t' \geq T$.

We first show that $(W_t)_{t=0}^{\tau}$ is increasing. We use that $W_\tau \geq W_t$ for all $t \leq \tau$, to show that $W_{\tau-1} \geq W_t$ for all $t \leq \tau - 1$. Take any $t \leq \tau - 1$ and note that since $\lambda_t \leq \lambda_{\tau-1}$ we have

$$\begin{aligned}
\lambda_t (b + \delta V) + (1 - \lambda_t) \max \{\delta W_0 - c, W_{t+1}\} &\leq \lambda_{\tau-1} (b + \delta V) + (1 - \lambda_{\tau-1}) \max \{\delta W_0 - c, \delta W_{t+1}\} \\
&\leq \lambda_{\tau-1} (b + \delta V) + (1 - \lambda_{\tau-1}) \max \{\delta W_0 - c, \delta W_\tau\}.
\end{aligned}$$

Thus, for any s,

$$\begin{aligned}
W_t (s) &= \max \{s + \delta W_0 - c, \lambda_t (b + \delta V) + (1 - \lambda_t) \max \{\delta W_0 - c, \delta W_{t+1}\}\} \\
&\leq \max \{s + \delta W_0 - c, \lambda_{\tau-1} (b + \delta V) + (1 - \lambda_{\tau-1}) \max \{\delta W_0 - c, \delta W_\tau\}\} = W_{\tau-1} (s).
\end{aligned}$$

Integrating over s, we obtain $W_t \leq W_{\tau-1}$. The same argument can be repeated for $\tau - 1$, since now we have that $W_{\tau-1} \geq W_t$ for all $t \leq \tau - 1$.

We are left to show that $(W_t)_{t=\tau}^{T}$ is decreasing. If $\tau = T$ this follows trivially, so consider $\tau < T$. We know that $W_{\tau+1} < W_\tau$ by definition, i.e.,

$$\begin{aligned}
&\int \max \{s + \delta W_0 - c, \lambda_{\tau+1} (b + \delta V) + (1 - \lambda_{\tau+1}) \max \{\delta W_0 - c, \delta W_{\tau+2}\}\} \, \mathrm{d}F (s) \\
< &\int \max \{s + \delta W_0 - c, \lambda_\tau (b + \delta V) + (1 - \lambda_\tau) \max \{\delta W_0 - c, \delta W_{\tau+1}\}\} \, \mathrm{d}F (s).
\end{aligned}$$

This can only hold if

$$\begin{aligned}
\lambda_{\tau+1} (b + \delta V) + (1 - \lambda_{\tau+1}) \max \{\delta W_0 - c, \delta W_{\tau+2}\} &< \lambda_\tau (b + \delta V) + (1 - \lambda_\tau) \max \{\delta W_0 - c, \delta W_{\tau+1}\} \\
&\leq \lambda_{\tau+1} (b + \delta V) + (1 - \lambda_{\tau+1}) \max \{\delta W_0 - c, \delta W_{\tau+1}\},
\end{aligned}$$

where the last line follows since $\lambda_\tau \leq \lambda_{\tau+1}$ and since $V \geq W_r$ for all r. The above implies $\max \{\delta W_0 - c, \delta W_{\tau+2}\} < \max \{\delta W_0 - c, \delta W_{\tau+1}\} = \delta W_{\tau+1}$. The equality follows from the fact that $\delta W_0 - c < \delta W_{\tau+1}$, as otherwise the inequality would fail. Thus $\max \{\delta W_0 - c, \delta W_{\tau+2}\} < \delta W_{\tau+1}$ and in particular $W_{\tau+2} < W_{\tau+1}$. The argument can be repeated to show that $W_{\tau+3} < W_{\tau+2}$, and so on, up to period T, since this is where the λ sequence is no longer increasing. ■

B.3 Proof of Theorem 3

Theorem. *There exists a $K \geq T$ such that $\delta W_t \geq \delta W_0 - c$ for all $t \leq K$ and $\delta W_t < \delta W_0 - c$ for all $t > K$.*

Proof. To show that the K defined in the statement satisfies $K \geq T$, we need to prove that $\delta W_T \geq \delta W_0 - c$.

Assume by way of contradiction that in the equilibrium, $\delta W_T < \delta W_0 - c$. Then by Proposition 2, $\delta W_{T+1} \leq \delta W_T < \delta W_0 - c$ and hence

$$W_T (s) = \max \{s + \delta W_0 - c, \lambda_T (b + \delta V) + (1 - \lambda_T) (\delta W_0 - c)\}.$$

Oberve that in that case, we must have

$$W_{T-1} (s) = \max \{s + \delta W_0 - c, \lambda_{T-1} (b + \delta V) + (1 - \lambda_{T-1}) (\delta W_0 - c)\}.$$

Note that since $\lambda_{T-1} \leq \lambda_T$ we have that $W_{T-1}(s) \leq W_T(s)$ for each s. Thus $W_{T-1} \leq W_T$ and hence $\delta W_{T-1} < \delta W_0 - c$.

We proceed by backward induction. Consider some $t \leq T-1$ such that $W_t < \delta W_0 - c$. Then we have that

$$\begin{aligned} W_{t-1}(s) &= \max\{s+\delta W_0 - c, \lambda_{t-1}(b+\delta V) + (1-\lambda_{t-1})(\delta W_0 - c)\} \\ &\leq \max\{s+\delta W_0 - c, \lambda_t(b+\delta V) + (1-\lambda_t)(\delta W_0 - c)\} = W_t(s). \end{aligned}$$

Iterating this until $t=1$, this implies $\delta W_0 \leq \delta W_0 - c$, which is a contradiction.

Proposition 2 implies that there exists a $K \geq T$ such that $\delta W_t \geq \delta W_0 - c$ for all $T \leq t \leq K$. We now want to show that $\delta W_t > \delta W_0 - c$ for all $t \leq T$. By Proposition 2, we have that $\min\{W_0, W_T\} \leq W_t$ for all $t \leq T$. But $\delta W_0 > \delta W_0 - c$ and $\delta W_T \geq \delta W_0 - c$, so that $\delta W_t \geq \delta W_0 - c$ for all $t \leq K$. □

Lemma 7. *We have that the total derivative* $\frac{dW_t}{dW_0} \leq \delta$ *for all* $t > 0$.

Proof. For any $t \geq K$, $W_t(s) = \max\{s+\delta W_0 - c, \lambda_t(b+\delta V) + (1-\lambda_t)(\delta W_0 - c)\}$. If $s \leq s_t^*$,

$$\frac{\mathrm{d}W_t(s)}{\mathrm{d}W_0} = \lambda_t \delta \frac{\mathrm{d}V}{\mathrm{d}W_0} + (1-\lambda_t)\delta \leq \lambda_t \delta^2 + (1-\lambda_t)\delta \leq \delta,$$

where the inequality follows from Lemma 1. If $s > s_t^*$, we have $\frac{\mathrm{d}W_t(s)}{\mathrm{d}W_0} = \delta$ and hence $\frac{\mathrm{d}W_t(s)}{\mathrm{d}W_0} \leq \delta$ for all s. Integrate over s and apply the Leibniz integral rule to get

$$\begin{aligned} \frac{\mathrm{d}W_t}{\mathrm{d}W_0} &= \frac{\mathrm{d}}{\mathrm{d}W_0}\int W_t(s, W_0)\ \mathrm{d}F(s) = \int_\ell^{s_t^*} \frac{\mathrm{d}W_t(s, W_0)}{\mathrm{d}W_0}\ \mathrm{d}F(s) + \int_{s_t^*}^\infty \frac{\mathrm{d}W_t(s, W_0)}{\mathrm{d}W_0}\ \mathrm{d}F(s) \\ &\leq F(s_t^*)\delta + (1-F(s_t^*))\delta = \delta. \end{aligned}$$

We complete the proof by backward induction. Given $\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} \leq \delta$, we will show that $\frac{\mathrm{d}W_t}{\mathrm{d}W_0} \leq \delta$ for all $t < K$. Since $t < K$, $W_t(s) = \max\{s+\delta W_0 - c, \lambda_t(b+\delta V) + (1-\lambda_t)\delta W_{t+1}\}$. If $s \leq s_t^*$,

$$\frac{\mathrm{d}W_t(s)}{\mathrm{d}W_0} = \lambda_t \delta \frac{\mathrm{d}V}{\mathrm{d}W_0} + (1-\lambda_t)\delta \frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} \leq \lambda_t\delta^2 + (1-\lambda_t)\delta^2 = \delta^2 \leq \delta,$$

where the first inequality follows from Lemma 1. If $s > s_t^*$, we have $\frac{\mathrm{d}W_t(s)}{\mathrm{d}W_0} = \delta$ and thus $\frac{\mathrm{d}W_t(s)}{\mathrm{d}W_0} \leq \delta$ for all s. Again integrating over s yields $\frac{\mathrm{d}W_t}{\mathrm{d}W_0} \leq \delta$. □

B.4 Proof of Proposition 2

Proposition. $W_{t'}$ *is increasing in* λ_t *for any* $t', t \leq K$. *Furthermore,* W_0 *is strictly increasing in* λ_t *if and only if* $F(s_k^*) > 0$, *for all* $0 \leq k \leq t$ *and* $\lambda_k < 1$ *for all* $0 < k < t-1$.

Proof. Fix $t \leq K$. For any $r \leq t$ define

$$\begin{aligned} \alpha_r &= (1-F(s_r^*))\delta + F(s_r^*)\lambda_r\delta\frac{\mathrm{d}V}{\mathrm{d}W_0}, \\ \beta_r &= F(s_r^*)(1-\lambda_r)\delta, \\ \gamma_r &= \alpha_r + \beta_r(\gamma_{r+1}), \end{aligned}$$

where we set $\gamma_{t+1} = \frac{dW_{t+1}}{dW_0}$ and define $\frac{dW_{K+1}}{dW_0} = 1$. Although these coefficients are functions of t, we drop this dependence to simplify notation. The difference in the definition of γ_{t+1} as opposed to γ_{K+1} accounts for the different definitions of the value function in those periods. It is immediate that for all $r \leq t$, $0 \leq \gamma_r \leq \delta$.

Recall that $W_K(s) = \max\{s + \delta W_0 - c, \lambda_K(b + \delta V) + (1 - \lambda_K)(\delta W_0 - c)\}$. If $s \leq s_K^*$,

$$\begin{aligned}\frac{\mathrm{d}W_K(s)}{\mathrm{d}\lambda_K} &= b + \delta V + \lambda_K \delta \frac{\mathrm{d}V}{\mathrm{d}W_0}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_K} + (1-\lambda_K)\delta\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_K} - \delta W_0 + c \\ &= \psi_K + \lambda_K \delta \frac{\mathrm{d}V}{\mathrm{d}W_0}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_K} + (1-\lambda_K)\delta\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_K},\end{aligned}$$

where we denote $\psi_K = b + \delta V - \delta W_0 + c$. For $s > s_K^*$, we have that $\frac{\mathrm{d}W_K(s)}{\mathrm{d}\lambda_K} = \delta\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_K}$.

To find $\frac{\mathrm{d}W_K}{\mathrm{d}\lambda_K}$ we need to integrate the above expressions with respect to s and apply the Leibniz integral rule (this applies since W_K is integrable and partial derivatives are bounded and exist almost everywhere). Doing this and simplifying we obtain

$$\begin{aligned}\frac{\mathrm{d}W_K}{\mathrm{d}\lambda_K} &= \int_\ell^{s_K^*} \frac{\mathrm{d}W_K(s,\lambda_K)}{\mathrm{d}\lambda_K}\,\mathrm{d}F(s) + \int_{s_K^*}^{\infty} \frac{\mathrm{d}W_K(s,\lambda_K)}{\mathrm{d}\lambda_K}\,\mathrm{d}F(s) \\ &= F(s_K^*)\left(\psi_K + \lambda_K\delta\frac{\mathrm{d}V}{\mathrm{d}W_0}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_K} + (1-\lambda_K)\delta\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_K}\right) + (1 - F(s_K^*))\delta\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_K} \\ &= F(s_K^*)\psi_K + \gamma_K\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_K},\end{aligned}$$

where $\gamma_K = \alpha_K + \beta_K$ as defined above (recall that $\gamma_{K+1} = 1$).

Now, for $t < K$, $W_t(s) = \max\{s + \delta W_0 - c, \lambda_t(b + \delta V) + (1-\lambda_t)\delta W_{t+1}\}$. For $s \leq s_t^*$,

$$\frac{\mathrm{d}W_t(s)}{\mathrm{d}\lambda_t} = \psi_t + \lambda_t\delta\frac{\mathrm{d}V}{\mathrm{d}W_0}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + (1-\lambda_t)\delta\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t},$$

where we define $\psi_t = b + \delta V - \delta W_{t+1}$. For $s > s_t^*$, $\frac{\mathrm{d}W_t(s)}{\mathrm{d}\lambda_t} = \delta\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}$. Integrating over s, we get

$$\begin{aligned}\frac{\mathrm{d}W_t}{\mathrm{d}\lambda_t} &= F(s_t^*)\left(\psi_t + \lambda_t\delta\frac{\mathrm{d}V}{\mathrm{d}W_0}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + (1-\lambda_t)\delta\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}\right) + (1 - F(s_t^*))\delta\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \\ &= F(s_t^*)\psi_t + \left(\lambda_t\delta F(s_t^*)\frac{\mathrm{d}V}{\mathrm{d}W_0} + (1-\lambda_t)\delta F(s_t^*)\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} + (1-F(s_t^*))\delta\right)\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \\ &= F(s_t^*)\psi_t + \gamma_t\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}. \qquad (18)\end{aligned}$$

The definition of γ_t implies that we have the same expression for $\frac{\mathrm{d}W_t}{\mathrm{d}\lambda_t}$ and $\frac{\mathrm{d}W_K}{\mathrm{d}\lambda_K}$ and can now thus consider any generic period $r < t$ where $W_r(s) = \max\{s + \delta W_0 - c, \lambda_r(b + \delta V) + (1-\lambda_r)\delta W_{r+1}\}$. Through a similar argument we find that

$$\begin{aligned}\frac{\mathrm{d}W_r}{\mathrm{d}\lambda_t} &= F(s_r^*)\left(\lambda_r\delta\frac{\mathrm{d}V}{\mathrm{d}W_0}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + (1-\lambda_r)\delta\frac{\mathrm{d}W_{r+1}}{\mathrm{d}\lambda_t}\right) + (1 - F(s_r^*))\delta\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \\ &= \left((1 - F(s_r^*))\delta + F(s_r^*)\lambda_r\delta\frac{\mathrm{d}V}{\mathrm{d}W_0}\right)\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + F(s_r^*)(1-\lambda_r)\delta\frac{\mathrm{d}W_{r+1}}{\mathrm{d}\lambda_t} \\ &= \alpha_r\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + \beta_r\frac{\mathrm{d}W_{r+1}}{\mathrm{d}\lambda_t}. \qquad (19)\end{aligned}$$

Working backwards from t to $r = t-1$, we have that

$$\frac{\mathrm{d}W_{t-1}}{\mathrm{d}\lambda_t}\alpha_{t-1}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + \beta_{t-1}\frac{\mathrm{d}W_t}{\mathrm{d}\lambda_t} = \gamma_{t-1}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + \beta_{t-1}F\left(s_t^*\right)\psi_t.$$

More generally, we get $\frac{\mathrm{d}W_r}{\mathrm{d}\lambda_t} = \gamma_r \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + \left(\prod_{k=r}^{t-1}\beta_k\right)F\left(s_t^*\right)\psi_t$. Letting $r = 0$, we get

$$\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} = \frac{\left(\prod_{k=r}^{t-1}\beta_k\right)F\left(s_t^*\right)\psi_t}{1-\gamma_0} = \frac{\left(\prod_{k=0}^{t-1}F\left(s_k^*\right)\left(1-\lambda_k\right)\right)F\left(s_t^*\right)\delta^t\psi_t}{1-\gamma_0} \geq 0, \tag{20}$$

where the inequality holds because $\gamma_0 \leq \delta$. Since the denominator is strictly positive, W_0 is strictly increasing in λ_t if $F\left(s_k^*\right) > 0$ for all $0 \leq k \leq t$ and $\lambda_k < 1$ for all $0 < k < t-1$, since that makes the numerator strictly positive.

Now, equation (18) along with the fact that $\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \geq 0$ shows that $\frac{\mathrm{d}W_t}{\mathrm{d}\lambda_t} > 0$, for all $t \leq K$. Furthermore, by backward induction from period t, equation (19) shows that $\frac{\mathrm{d}W_\tau}{\mathrm{d}\lambda_t} \geq 0$ for all $\tau < t$. Note that this is strict as long as $F\left(s_r^*\right) > 0$ and $\lambda_r < 1$ for all r with $\tau \leq r < t$.

Finally for $\tau > t$, $\frac{\mathrm{d}W_\tau}{\mathrm{d}\lambda_t} = \frac{\mathrm{d}W_\tau}{\mathrm{d}W_0}\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}$, we have shown above that $\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \geq 0$ and lemma 7 shows that $\frac{\mathrm{d}W_\tau}{\mathrm{d}W_0} \geq 0$ (with the inequality strict if $\lambda_\tau < 1$). □

B.5 Proof of Proposition 3

Proposition. *For each $t \leq K$, the total derivative $\frac{\mathrm{d}s_t^*}{\mathrm{d}W_0} < 0$. Furthermore $\frac{\mathrm{d}s^*}{\mathrm{d}W_0} < 0$.*

Proof. By equation (6), $s_K^* = \lambda_K\left(b + \delta V - \delta W_0 + c\right)$. Differentiating with respect to W_0 yields $\frac{\mathrm{d}s_K^*}{\mathrm{d}W_0} = \lambda_K\delta\left(\frac{\mathrm{d}V}{\mathrm{d}W_0} - 1\right) \leq 0$, since $\frac{\mathrm{d}V}{\mathrm{d}W_0} = \frac{1-F(s^*)}{1/\delta - F(s^*)} \in [0,\delta]$. Similarly, by differentiating equation (3) with respect to W_0 we get $\frac{\mathrm{d}s^*}{\mathrm{d}W_0} = \delta\left(\frac{\mathrm{d}V}{\mathrm{d}W_0} - 1\right) \leq 0$.

From (6), $s_t^* = \lambda_t\left(b+\delta V\right) + \left(1-\lambda_t\right)\delta W_{t+1} - \delta W_0 + c$, for any $t < K$. So

$$\frac{\mathrm{d}s_t^*}{\mathrm{d}W_0} = \lambda_t\delta\frac{\mathrm{d}V}{\mathrm{d}W_0} + \left(1-\lambda_t\right)\delta\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} - \delta = \lambda_t\delta\left(\frac{\mathrm{d}V}{\mathrm{d}W_0} - 1\right) + \left(1-\lambda_t\right)\delta\left(\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} - 1\right) \leq -\delta\left(1-\delta\right) < 0,$$

since $\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} \leq \delta$ by Lemma 7 and $\frac{\mathrm{d}V}{\mathrm{d}W_0} \leq \delta$ by Lemma 1. □

B.6 Proof of Proposition 4

Proposition. *For each $t \leq K$, $\frac{\mathrm{d}s_\tau^*}{\mathrm{d}\lambda_t} \geq 0$ if $\tau \leq t$ and $\frac{\mathrm{d}s_\tau^*}{\mathrm{d}\lambda_t} \leq 0$ if $\tau > t$.*

Proof. Fix $t < \tau < K$. By equation (6), $s_\tau^* = \lambda_\tau\left(b+\delta V\right) + \left(1-\lambda_\tau\right)\delta W_{\tau+1} - \delta W_0 + c$ and

$$\begin{aligned}
\frac{\mathrm{d}s_\tau^*}{\mathrm{d}\lambda_t} &= \lambda_\tau\delta\left(\frac{\mathrm{d}V}{\mathrm{d}\lambda_t} - \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}\right) + \left(1-\lambda_\tau\right)\delta\left(\frac{\mathrm{d}W_{\tau+1}}{\mathrm{d}\lambda_t} - \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}\right) \\
&= \lambda_\tau\delta\left(\frac{\mathrm{d}V}{\mathrm{d}W_0} - 1\right)\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + \left(1-\lambda_\tau\right)\delta\left(\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} - 1\right)\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \\
&\leq \lambda_\tau\delta\left(\delta-1\right)\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + \left(1-\lambda_\tau\right)\delta\left(\delta-1\right)\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} = \delta\left(\delta-1\right)\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \leq 0,
\end{aligned}$$

where the third line follows from Lemma 7.

Consider $t < \tau = K$. By equation (6), $s_K^* = \lambda_K (b + \delta V - \delta W_0 + c)$ and

$$\frac{\mathrm{d}s_K^*}{\mathrm{d}\lambda_t} = \lambda_K \delta \left(\frac{\mathrm{d}V}{\mathrm{d}\lambda_t} - \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \right) = \lambda_K \delta \left(\frac{\mathrm{d}V}{\mathrm{d}W_0} - 1 \right) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \leq 0,$$

since $\frac{\mathrm{d}V}{\mathrm{d}W_0} \leq \delta$ by Lemma 1 and $\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \geq 0$ by Proposition 2.

By equation (6), for any $t < K$ we have $s_t^* = \lambda_t (b + \delta V) + (1 - \lambda_t) \delta W_{t+1} - \delta W_0 + c$, so that

$$\begin{aligned} \frac{\mathrm{d}s_t^*}{\mathrm{d}\lambda_t} &= b + \delta V - \delta W_{t+1} + \lambda_t \delta \left(\frac{\mathrm{d}V}{\mathrm{d}\lambda_t} - \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \right) + (1 - \lambda_t) \delta \left(\frac{\mathrm{d}W_{t+1}}{\mathrm{d}\lambda_t} - \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \right) \\ &= b + \delta V - \delta W_{t+1} + \lambda_t \delta \left(\frac{\mathrm{d}V}{\mathrm{d}W_0} - 1 \right) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + (1 - \lambda_t) \delta \left(\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} - 1 \right) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \\ &= \psi_t + \lambda_t \delta \left(\frac{\mathrm{d}V}{\mathrm{d}W_0} - 1 \right) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + (1 - \lambda_t) \delta \left(\frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} - 1 \right) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}. \end{aligned}$$

We have that $\frac{\mathrm{d}s_t^*}{\mathrm{d}\lambda_t} \geq 0$ if and only if

$$\begin{aligned} \psi_t &\geq \delta \left[\lambda_t \left(1 - \frac{\mathrm{d}V}{\mathrm{d}W_0} \right) + (1 - \lambda_t) \left(1 - \frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} \right) \right] \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \\ &= \delta \left(1 - \lambda_t \frac{\mathrm{d}V}{\mathrm{d}W_0} - (1 - \lambda_t) \frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} \right) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}. \end{aligned}$$

Since $\delta - \gamma_t = \delta F(s_t^*) \left(1 - \lambda_t \frac{\mathrm{d}V}{\mathrm{d}W_0} - (1 - \lambda_t) \frac{\mathrm{d}W_{t+1}}{\mathrm{d}W_0} \right)$, we have $\frac{\mathrm{d}s_t^*}{\mathrm{d}\lambda_t} \geq 0$ if and only if $\psi_t \geq \frac{\delta - \gamma_t}{F(s_t^*)} \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}$, or $\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \leq \frac{F(s_t^*)\psi_t}{\delta - \gamma_t}$. We shall show something stronger, namely that

$$\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} = \frac{\left(\prod_{k=r}^{t-1} \beta_k \right) F(s_t^*) \psi_t}{1 - \gamma_0} \leq \frac{F(s_t^*) \psi_t}{1 - \gamma_t}.$$

Observe that $t = 0$ holds trivially. For arbitrary t, it suffices to show that $\left(\prod_{k=r}^{t-1} \beta_k \right) (1 - \gamma_t) \leq 1 - \gamma_0$. To prove this we show that for all $r \leq t$, $1 - \gamma_r \geq \beta_r (1 - \gamma_{r+1})$. This follows since

$$\begin{aligned} 1 - \gamma_r &= 1 - \alpha_r - \beta_r (\gamma_{r+1}) \\ &= 1 - (1 - F(s_r^*)) \delta - F(s_r^*) \lambda_r \delta \frac{\mathrm{d}V}{\mathrm{d}W_0} - \beta_r (\gamma_{r+1}) \\ &\geq 1 - \delta + \delta F(s_r^*) - F(s_r^*) \lambda_r \delta - \beta_r (\gamma_{r+1}) \\ &\geq \delta F(s_r^*) - F(s_r^*) \lambda_r \delta - \beta_r (\gamma_{r+1}) = \beta_r - \beta_r (\gamma_{r+1}). \end{aligned}$$

The case of $t = K$ is proved mutatis mutandis. Thus we have shown that $\frac{\mathrm{d}s_t^*}{\mathrm{d}\lambda_t} \geq 0$ for all t.

We are left to prove that $\frac{\mathrm{d}s_\tau^*}{\mathrm{d}\lambda_t} = \lambda_\tau \delta \left(\frac{\mathrm{d}V}{\mathrm{d}W_0} - 1 \right) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + (1 - \lambda_\tau) \delta \left(\frac{\mathrm{d}W_{\tau+1}}{\mathrm{d}\lambda_t} - \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \right) \geq 0$, for any $\tau < t \leq K$. This is equivalent to showing $\beta_\tau \frac{\mathrm{d}W_{\tau+1}}{\mathrm{d}\lambda_t} \geq \left(F(s_\tau^*) \delta - F(s_\tau^*) \delta \lambda_\tau \frac{\mathrm{d}V}{\mathrm{d}W_0} \right) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}$. And, since $1 - \alpha_\tau \geq F(s_\tau^*) \delta - F(s_\tau^*) \delta \lambda_\tau \frac{\mathrm{d}V}{\mathrm{d}W_0}$, it suffices to show that

$$\beta_\tau \frac{\mathrm{d}W_{\tau+1}}{\mathrm{d}\lambda_t} \geq (1 - \alpha_\tau) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}. \tag{21}$$

Consider $\tau = 0 < t \leq K$. Since $\frac{\mathrm{d}W_r}{\mathrm{d}\lambda_t} = \alpha_r \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + \beta_r \frac{\mathrm{d}W_{r+1}}{\mathrm{d}\lambda_t}$, setting $r = 0$ we have $\beta_0 \frac{\mathrm{d}W_1}{\mathrm{d}\lambda_t} = (1 - \alpha_0) \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_1}$, so the above holds trivially. Consider any $0 < \tau < t \leq K$. Setting $r = \tau$, we have

$\beta_\tau \frac{\mathrm{d}W_{\tau+1}}{\mathrm{d}\lambda_t} = \frac{\mathrm{d}W_\tau}{\mathrm{d}\lambda_t} - \alpha_\tau \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}$. So, to prove (21), we need to show that $\frac{\mathrm{d}W_\tau}{\mathrm{d}\lambda_t} \geq \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}$.

We proceed by induction on τ. The case $\tau = 0$ holds since $\frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} = \alpha_0 \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + \beta_0 \frac{\mathrm{d}W_1}{\mathrm{d}\lambda_t}$, and so

$$\frac{\mathrm{d}W_1}{\mathrm{d}\lambda_t} = \frac{1-\alpha_0}{\beta_0} \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \geq \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t},$$

since $1-\alpha_0 \geq \beta_0$, because $\alpha_0 + \beta_0 \leq \delta < 1$. For the inductive step, assume that $\frac{\mathrm{d}W_{\tau-1}}{\mathrm{d}\lambda_t} \geq \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}$. We claim that $\frac{\mathrm{d}W_\tau}{\mathrm{d}\lambda_t} \geq \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t}$. Because $\frac{\mathrm{d}W_r}{\mathrm{d}\lambda_t} = \alpha_r \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} + \beta_r \frac{\mathrm{d}W_{r+1}}{\mathrm{d}\lambda_t}$, we have

$$\frac{\mathrm{d}W_\tau}{\mathrm{d}\lambda_t} = \frac{1}{\beta_{\tau-1}} \frac{\mathrm{d}W_{\tau-1}}{\mathrm{d}\lambda_t} - \frac{\alpha_{\tau-1}}{\beta_{\tau-1}} \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \geq \frac{1}{\beta_{\tau-1}} \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} - \frac{\alpha_{\tau-1}}{\beta_{\tau-1}} \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} = \frac{1-\alpha_{\tau-1}}{\beta_{\tau-1}} \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t} \geq \frac{\mathrm{d}W_0}{\mathrm{d}\lambda_t},$$

where the last inequality follows since $1-\alpha_{\tau-1} \geq \beta_{\tau-1}$, because $\alpha_{\tau-1} + \beta_{\tau-1} \leq \delta < 1$. □

B.7 Proof of Proposition 5

Proposition. *The agent's W_t is increasing in δ, i.e., $\frac{\mathrm{d}W_t}{\mathrm{d}\delta} > 0$ for all $t \leq K$.*

Proof. To prove the Proposition, we start with the following claim.

Claim: The total derivative $\frac{\mathrm{d}V}{\mathrm{d}\delta} = \alpha + \beta \frac{\mathrm{d}W_0}{\mathrm{d}\delta}$, where $\alpha > W_0$ and $\delta > \beta \geq 0$.

Proof of Claim: By the definition of V in equation (2), we have

$$V = \int_\ell^{s^*} (b + \delta V) \ \mathrm{d}F(s) + \int_{s^*}^\infty (s + \delta W_0 - c) \ \mathrm{d}F(s).$$

The total derivative of both sides with respect to δ is

$$\frac{\mathrm{d}V}{\mathrm{d}\delta} = F(s^*) \left(V + \delta \frac{\mathrm{d}V}{\mathrm{d}\delta} \right) + (1 - F(s^*)) \left(W_0 + \delta \frac{\mathrm{d}W_0}{\mathrm{d}\delta} \right),$$

and solving for $\frac{\mathrm{d}V}{\mathrm{d}\delta}$ we find

$$\frac{\mathrm{d}V}{\mathrm{d}\delta} = \frac{F(s^*) V + (1 - F(s^*)) W_0}{1 - \delta F(s^*)} + \frac{\delta - \delta F(s^*)}{1 - \delta F(s^*)} \frac{\mathrm{d}W_0}{\mathrm{d}\delta} = \alpha + \beta \frac{\mathrm{d}W_0}{\mathrm{d}\delta},$$

where $\alpha = \frac{F(s^*)V + (1-F(s^*))W_0}{1-\delta F(s^*)} > W_0$ and $\delta > \beta = \frac{\delta - \delta F(s^*)}{1-\delta F(s^*)} \geq 0$ where the first inequality follows since $F(s^*) > 0$ in any relational contracting equilibrium. This proves the claim. ■

Observe that

$$\frac{\alpha}{1-\beta} = \frac{F(s^*) V + (1 - F(s^*)) W_0}{1 - \delta}.$$

With α and β defined above, note that $\alpha = F(s^*)(V + \delta\alpha) + (1 - F(s^*)) W_0$ and $\beta = F(s^*) \delta\beta + (1 - F(s^*)) \delta$, so that $1 - \beta = 1 - \delta + F(s^*) \delta (1 - \beta)$.

If $t = K$, equation (5) implies $W_K(s) = \max\{s + \delta W_0 - c, \lambda_K (b + \delta V) + (1 - \lambda_K)(\delta W_0 - c)\}$. So that if $s > s_K^*$, $\frac{\mathrm{d}W_K(s)}{\mathrm{d}\delta} = W_0 + \delta \frac{\mathrm{d}W_0}{\mathrm{d}\delta}$ and if $s \leq s_K^*$ then

$$\frac{\mathrm{d}W_K(s)}{\mathrm{d}\delta} = \lambda_K \left(V + \delta \frac{\mathrm{d}V}{\mathrm{d}\delta} \right) + (1 - \lambda_K) \left(W_0 + \delta \frac{\mathrm{d}W_0}{\mathrm{d}\delta} \right).$$

Integrating over s, as usual, and simplifying we get

$$\begin{aligned}\frac{\mathrm{d}W_K}{\mathrm{d}\delta} &= F\left(s_K^*\right)\left(\lambda_K\left(V+\delta\frac{\mathrm{d}V}{\mathrm{d}\delta}\right)+\left(1-\lambda_K\right)\left(W_0+\delta\frac{\mathrm{d}W_0}{\mathrm{d}\delta}\right)\right)+\left(1-F\left(s_K^*\right)\right)\left(W_0+\delta\frac{\mathrm{d}W_0}{\mathrm{d}\delta}\right)\\ &= W_0+F\left(s_K^*\right)\lambda_K\left(V+\delta\alpha-W_0\right)+\left[\delta-\delta F\left(s_K^*\right)\lambda_K\left(1-\beta\right)\right]\frac{\mathrm{d}W_0}{\mathrm{d}\delta}=\alpha_K+\beta_K\frac{\mathrm{d}W_0}{\mathrm{d}\delta},\end{aligned}$$

where we now define $\alpha_K = F\left(s_K^*\right)\lambda_K V + \left(1-F\left(s_K^*\right)\lambda_K\right)W_0 + F\left(s_K^*\right)\lambda_K\delta\alpha$ and $\beta_K = \delta - \delta F\left(s_K^*\right)\lambda_K\left(1-\beta\right)$. Note that

$$\frac{\alpha}{1-\beta}=\frac{F\left(s^*\right)V+\left(1-F\left(s^*\right)\right)W_0}{1-\delta}\geq\frac{F\left(s_K^*\right)\lambda_K V+\left(1-F\left(s_K^*\right)\lambda_K\right)W_0}{1-\delta},$$

since $F\left(s^*\right)\geq F\left(s_K^*\right)$ and $\lambda_K\leq 1$. Letting $\xi=F\left(s_K^*\right)\lambda_K$, the above then implies

$$\begin{aligned}\alpha\left(1-\delta\right) &\geq \left(1-\beta\right)\left(\xi_K V+\left(1-\xi_K\right)W_0\right)\\ \alpha\left(1-\delta\right)+\delta\xi_K\alpha\left(1-\beta\right) &\geq \left(1-\beta\right)\left(\xi_K V+\left(1-\xi_K\right)W_0\right)+\delta\xi_K\alpha\left(1-\beta\right)\\ \alpha(1-\delta+\delta\xi_K\left(1-\beta\right)) &\geq \left(1-\beta\right)\left(\xi_K V+\left(1-\xi_K\right)W_0+\delta\xi_K\alpha\right)\\ \frac{\alpha}{1-\beta} &\geq \frac{\xi_K V+\left(1-\xi_K\right)W_0+\delta\xi_K\alpha}{1-\delta+\delta\xi_K\left(1-\beta\right)}\geq\frac{\alpha_K}{1-\beta_K}.\end{aligned}$$

Now, $\beta_K \geq \beta$ since $\beta_K = F\left(s_K^*\right)\lambda_K\beta + \left(1-F\left(s_K^*\right)\lambda_K\right)\delta \geq F\left(s^*\right)\delta\beta + \left(1-F\left(s^*\right)\right)\delta = \beta$, where the inequality follows because $\beta<\delta$, $F\left(s^*\right)\geq F\left(s_K^*\right)$ and $\lambda_K\leq 1$. Finally, $\alpha_K\leq\alpha$ since

$$\begin{aligned}\alpha_K &= W_0+F\left(s_K^*\right)\lambda_K\left(V+\delta\alpha-W_0\right)\\ &= \delta F\left(s_K^*\right)\lambda_K\alpha+F\left(s_K^*\right)\lambda_K V+\left(1-F\left(s_K^*\right)\lambda_K\right)W_0\\ &\leq \delta F\left(s^*\right)\alpha+F\left(s^*\right)V+\left(1-F\left(s^*\right)\right)W_0=\alpha.\end{aligned}$$

For $t<K$, by equation (5) we have $W_t\left(s\right)=\max\left\{s+\delta W_0-c, \lambda_t\left(b+\delta V\right)+\left(1-\lambda_t\right)\delta W_{t+1}\right\}$. So that if $s>s_t^*$, $\frac{\mathrm{d}W_t(s)}{\mathrm{d}\delta}=W_0+\delta\frac{\mathrm{d}W_0}{\mathrm{d}\delta}$ and if $s\leq s_t^*$ then

$$\frac{\mathrm{d}W_t\left(s\right)}{\mathrm{d}\delta}=\lambda_t\left(V+\delta\frac{\mathrm{d}V}{\mathrm{d}\delta}\right)+\left(1-\lambda_t\right)\left(W_{t+1}+\delta\frac{\mathrm{d}W_{t+1}}{\mathrm{d}\delta}\right).$$

Integrating over s, and simplifying results in

$$\begin{aligned}\frac{\mathrm{d}W_t}{\mathrm{d}\delta} &= F\left(s_t^*\right)\left(\lambda_t\left(V+\delta\frac{\mathrm{d}V}{\mathrm{d}\delta}\right)+\left(1-\lambda_t\right)\left(W_{t+1}+\delta\frac{\mathrm{d}W_{t+1}\left(s\right)}{\mathrm{d}\delta}\right)\right)+\left(1-F\left(s_t^*\right)\right)\left(W_0+\delta\frac{\mathrm{d}W_0}{\mathrm{d}\delta}\right)\\ &= F\left(s_t^*\right)\left(\lambda_t V+\delta\lambda_t\alpha+\left(1-\lambda_t\right)W_{t+1}\right)+\left(1-F\left(s_t^*\right)\right)W_0\\ &\quad +F\left(s_t^*\right)\delta\lambda_t\beta\frac{\mathrm{d}W_0}{\mathrm{d}\delta}+F\left(s_t^*\right)\delta\left(1-\lambda_t\right)\frac{\mathrm{d}W_{t+1}\left(s\right)}{\mathrm{d}\delta}+\left(1-F\left(s_t^*\right)\right)\delta\frac{\mathrm{d}W_0}{\mathrm{d}\delta}.\end{aligned}$$

We proceed by backward induction. Assuming that $\frac{\mathrm{d}W_t}{\mathrm{d}\delta}=\alpha_t+\beta_t\frac{\mathrm{d}W_0}{\mathrm{d}\delta}$ where $\alpha\geq\alpha_t>0$

and $\delta \geq \beta_t \geq \beta$ we have that

$$\begin{aligned}
\frac{\mathrm{d}W_{t-1}}{\mathrm{d}\delta} &= F\left(s_{t-1}^*\right)\left(\lambda_{t-1}V + \delta\lambda_{t-1}\alpha + \left(1-\lambda_{t-1}\right)W_t\right) + \left(1-F\left(s_{t-1}^*\right)\right)W_0 \\
&\quad + \left[F\left(s_{t-1}^*\right)\delta\lambda_{t-1}\beta + \left(1-F\left(s_{t-1}^*\right)\right)\delta\right]\frac{\mathrm{d}W_0}{\mathrm{d}\delta} + F\left(s_{t-1}^*\right)\delta\left(1-\lambda_{t-1}\right)\frac{\mathrm{d}W_t\left(s\right)}{\mathrm{d}\delta} \\
&= F\left(s_{t-1}^*\right)\left(\lambda_{t-1}V + \delta\lambda_{t-1}\alpha + \left(1-\lambda_{t-1}\right)W_t\right) + \left(1-F\left(s_{t-1}^*\right)\right)W_0 + F\left(s_{t-1}^*\right)\delta\left(1-\lambda_{t-1}\right)\alpha_t \\
&\quad + \left[F\left(s_{t-1}^*\right)\delta\lambda_{t-1}\beta + \left(1-F\left(s_{t-1}^*\right)\right)\delta + F\left(s_{t-1}^*\right)\delta\left(1-\lambda_{t-1}\right)\beta_t\right]\frac{\mathrm{d}W_0}{\mathrm{d}\delta} \\
&= \alpha_{t-1} + \beta_{t-1}\frac{\mathrm{d}W_0}{\mathrm{d}\delta},
\end{aligned}$$

where $\alpha_{t-1} = F\left(s_{t-1}^*\right)\left(\lambda_{t-1}V + \delta\lambda_{t-1}\alpha + \left(1-\lambda_{t-1}\right)W_t\right) + \left(1-F\left(s_{t-1}^*\right)\right)W_0 + F\left(s_{t-1}^*\right)\delta\left(1-\lambda_{t-1}\right)\alpha_t > 0$ and $\beta_{t-1} = F\left(s_{t-1}^*\right)\delta\lambda_{t-1}\beta + \left(1-F\left(s_{t-1}^*\right)\right)\delta + F\left(s_{t-1}^*\right)\delta\left(1-\lambda_{t-1}\right)\beta_t$.

Note that

$$\frac{\alpha}{1-\beta} = \frac{F\left(s^*\right)V + \left(1-F\left(s^*\right)\right)W_0}{1-\delta} \geq \frac{F\left(s_{t-1}^*\right)V + \left(1-F\left(s_{t-1}^*\right)\right)W_0}{1-\delta},$$

since $F\left(s^*\right) \geq F\left(s_{t-1}^*\right)$. Thus $\alpha\left(1-\delta\right) \geq \left(1-\beta\right)\left(F\left(s_{t-1}^*\right)V + \left(1-F\left(s_{t-1}^*\right)\right)W_0\right)$ and

$$\begin{aligned}
\alpha\left(1-\delta+\delta F\left(s_{t-1}^*\right)\left(1-\beta_t\right)\right) &\geq \left(1-\beta\right)\left(F\left(s_{t-1}^*\right)V + \left(1-F\left(s_{t-1}^*\right)\right)W_0\right) + \left(1-\beta_t\right)F\left(s_{t-1}^*\right)\delta\alpha \\
&\geq \left(1-\beta\right)\left(F\left(s_{t-1}^*\right)V + \left(1-F\left(s_{t-1}^*\right)\right)W_0\right) + \left(1-\beta\right)F\left(s_{t-1}^*\right)\delta\alpha \\
&= \left(1-\beta\right)\left(F\left(s_{t-1}^*\right)\left(V+\delta\alpha\right) + \left(1-F\left(s_{t-1}^*\right)\right)W_0\right).
\end{aligned}$$

We can therefore write

$$\begin{aligned}
\frac{\alpha}{1-\beta} &\geq \frac{F\left(s_{t-1}^*\right)\left(V+\delta\alpha\right) + \left(1-F\left(s_{t-1}^*\right)\right)W_0}{1-\delta+\delta F\left(s_{t-1}^*\right)\left(1-\beta_t\right)} \qquad (22) \\
&\geq \frac{F\left(s_{t-1}^*\right)\left(\lambda_{t-1}\left(V+\delta\alpha\right) + \left(1-\lambda_{t-1}\right)\left(W_t+\delta\alpha_t\right)\right) + \left(1-F\left(s_{t-1}^*\right)\right)W_0}{1-\delta\left(1-F\left(s_{t-1}^*\right)\right) - \delta F\left(s_{t-1}^*\right)\beta_t} \\
&\geq \frac{F\left(s_{t-1}^*\right)\left(\lambda_{t-1}\left(V+\delta\alpha\right) + \left(1-\lambda_{t-1}\right)\left(W_t+\delta\alpha_t\right)\right) + \left(1-F\left(s_{t-1}^*\right)\right)W_0}{1-\delta\left(1-F\left(s_{t-1}^*\right)\right) - \delta F\left(s_{t-1}^*\right)\left(\lambda_t\beta + \left(1-\lambda_t\right)\beta_t\right)} = \frac{\alpha_{t-1}}{1-\beta_{t-1}}.
\end{aligned}$$

Thus $\alpha \geq \alpha_t > 0$, $\delta \geq \beta_t \geq \beta$ and $\frac{\alpha}{1-\beta} \geq \frac{\alpha_t}{1-\beta_t}$ for all $t \leq K$. In particular, at $t=0$ we get $\frac{\mathrm{d}W_0}{\mathrm{d}\delta} = \alpha_0 + \beta_0\frac{\mathrm{d}W_0}{\mathrm{d}\delta} = \frac{\alpha_0}{1-\beta_0} > 0$, since $\alpha_0 > 0$ and $\beta_0 \leq \delta$. Going back to $t=K$, we have that $\frac{\mathrm{d}W_K}{\mathrm{d}\delta} = \alpha_K + \beta_K\frac{\mathrm{d}W_0}{\mathrm{d}\delta} > 0$, since all terms are positive and $\alpha_K > 0$. Similarly, we have $\frac{\mathrm{d}W_t}{\mathrm{d}\delta} > 0$ for any $t \leq K$. Finally, $\frac{\mathrm{d}V}{\mathrm{d}\delta} = \alpha + \beta\frac{\mathrm{d}W_0}{\mathrm{d}\delta} > 0$, where the equality follows from the claim. □

B.8 Proof of Empirical Implication 4

Empirical Implication. *As the agent becomes more patient, he is less likely to end a productive relationship.*

Proof. We will show that the cutoff for terminating a productive relationship, s^*, is increasing in δ. Since $s^* = b + \delta V - \delta W_0 + c$, we have

$$\frac{\mathrm{d}s^*}{\mathrm{d}\delta} = V - W_0 + \delta\left(\frac{\mathrm{d}V}{\mathrm{d}\delta} - \frac{\mathrm{d}W_0}{\mathrm{d}\delta}\right).$$

Now, from the proof of Proposition 5 we get

$$\frac{\mathrm{d}V}{\mathrm{d}\delta} - \frac{\mathrm{d}W_0}{\mathrm{d}\delta} = \alpha + \beta\frac{\mathrm{d}W_0}{\mathrm{d}\delta} - \frac{\mathrm{d}W_0}{\mathrm{d}\delta} = \alpha - (1-\beta)\frac{\mathrm{d}W_0}{\mathrm{d}\delta} = \alpha - (1-\beta)\frac{\alpha_0}{1-\beta_0} \geq 0,$$

where the inequality follows from (22). Since $V - W_0 \geq 0$, $\frac{\mathrm{d}s^*}{\mathrm{d}\delta} \geq 0$ and strictly so if $V > W_0$. □

ONLINE APPENDIX

Building and Maintaining Productive Relationships

by Nemanja Antić, Ameet Morjaria, and Miguel Ángel Talamas Marcos

Table of Contents

A Principal Incentives and MPE

Recall that for simplicity we are assuming that $\lambda_0 > 0$ and that $\lambda_k = 0$ for all $k \geq 1$. This means that $T = K = 0$, $s^* = b + \delta V - \delta W_0 + c$ and $s_0^* = \lambda_0 s^*$. From equation (5), we have

$$\begin{aligned} W_0 &= F(s_0^*)(\lambda_0(b+\delta V) + (1-\lambda_0)(\delta W_0 - c)) + \int_{s_0^*}^{\infty} s + \delta W_0 - c \, dF(s) \\ &= F(s_0^*)(s_0^* + \delta W_0 - c) + \int_{s_0^*}^{\infty} s + \delta W_0 - c \, \mathrm{d}F(s) \\ &= \int_{\ell}^{\infty} \max\{s_0^*, s\} \, \mathrm{d}F(s) + \delta W_0 - c = \frac{1}{1-\delta}\left(\int_{\ell}^{\infty} \max\{\lambda_0 s^*, s\} \, \mathrm{d}F(s) - c\right), \end{aligned}$$

which is only a function of s^*. Observe that the agent's expected utility from a new relationship, W_0, is increasing in s^*. Similarly, we can rearrange equation (2) to get

$$\begin{aligned} V &= \int_{\ell}^{\infty} \max\{b + \delta V, s + \delta W_0 - c\} \, \mathrm{d}F(s) = \int_{\ell}^{\infty} \max\{b + \delta V, s - s^* + b + \delta V\} \, \mathrm{d}F(s) \\ &= \frac{1}{1-\delta}\left(b + \int_{\ell}^{\infty} \max\{0, s - s^*\} \, \mathrm{d}F(s)\right). \end{aligned}$$

Substituting W_0 and V into the definition of s^* we get

$$\begin{aligned} (1-\delta) s^* &= b + c + \delta \int_{\ell}^{\infty} \max\{0, s - s^*\} \, \mathrm{d}F(s) - \delta \int_{\ell}^{\infty} \max\{\lambda_0 s^*, s\} \, \mathrm{d}F(s) \\ s^* &= b + c + \delta \int_{\ell}^{\infty} \max\{s^*, s\} \, \mathrm{d}F(s) - \delta \int_{\ell}^{\infty} \max\{\lambda_0 s^*, s\} \, \mathrm{d}F(s) \\ &= b + c + \delta \int_{\lambda_0 s^*}^{s^*} (s^* - s) \, dF(s) + \delta (s^* - \lambda_0 s^*) F(\lambda_0 s^*). \end{aligned}$$

Integration by parts yields $\int_{\lambda_0 s^*}^{s^*} (s^* - s)\, dF(s) = [(s^* - s) F(s)]_{\lambda_0 s^*}^{s^*} - \int_{\lambda_0 s^*}^{s^*} F(s)\, d(s^* - s)$ and after simplifying we get $s^* = b + c + \delta \int_{\lambda_0 s^*}^{s^*} F(s)\, ds$.

Recall that $R(\lambda_0) = \lambda_0 R_p(\lambda_0) + (1-\lambda_0)\delta\rho$. Her continuation value from a productive relationship is $R_p(\lambda_0) = F(s^*)(\varsigma - b + \delta R_p(\lambda_0)) + (1 - F(s^*))\delta\rho$ and the principal's IC constraint holds as long as $\varsigma - b + \delta R_p(\lambda_0) \geq \varsigma + \delta\rho$. In the agent-preferred MPE, $b = \delta R_p(\lambda_0) - \delta\rho = \delta F(s^*)\varsigma - \delta(1-\delta)\rho$. If $F(s^*) = 1$, the bonus $b = \delta\varsigma - \delta(1-\delta)\rho$, i.e., it is effectively exogenous. In the more general framework of this extension, we can find and we can substitute this expression for b into our expression for s^* to get

$$s^* = \delta\varsigma F(s^*) - \delta(1-\delta)\rho + c + \delta \int_{\lambda_0 s^*}^{s^*} F(s)\, ds. \tag{23}$$

If F is absolutely continuous, the above has a fixed point by the intermediate value theorem.

To show that the principal's continuation utility $R(\lambda_0)$ could be non-monotone in λ_0 in this model, we make some convenient parametric assumptions. First we specialize to the case where F is the uniform distribution on $[1, 3]$, so that we obtain a quadratic in s^*. Let $\varsigma = 4/\delta$, $c = 1$ and $\rho = 0$ to further simplify the coefficients. Equation (23) can now be written as

$$s^* = 1 - \frac{\delta(1-\lambda_0^2)}{4}(s^*)^2 + \frac{\delta}{2}(1-\lambda_0)s^*.$$

The positive solution for s^* is $s^* = 4 \,/\, \left(2 - \delta(1-\lambda_0) + \sqrt{4 + \delta(1-\lambda_0)(4\lambda_0 + \delta(1-\lambda_0))}\right)$, which is decreasing in λ_0, so that higher clarity types are more likely to end a productive relationship. The principal's value from starting a relationship with an agent when clarity is λ_0 is $R(\lambda_0) = \lambda_0 F(s^*)$. As $\lambda_0 \to 1$, $s^* \to 1$ and $F(s^*) \to 0$, we have that $R(\lambda_0) \to 0$. Furthermore, $R(0) = 0$, but $R(\lambda_0) > 0$ if $\lambda_0 \in (0,1)$. Thus the principal's payoff is non-monotone in λ_0. Smoother choices for the distribution of shocks make this non-monotonicity easier to obtain.

B Comparison to Jovanovic (1979)

We begin by summarizing the two-period Jovanovic (1979) model, using Section 6.8 of Ljungqvist and Sargent (2018).[37] The true distribution of worker-firm match productivity is $\theta \sim N\left(\mu, \sigma_0^2\right)$. Assume that match productivity, θ, is observed with noise in period 1, where the noise $u \sim N\left(0, \sigma_u^2\right)$. Match productivity is perfectly observed in period 2. The worker's wage is the expected match productivity in every period. In the first period, the firm offers the worker a wage of $m_0 \sim N\left(\mu, \sigma_1^2\right)$, where $\sigma_1^2 = K_0 \sigma_u^2$ and $K_0 = \frac{\sigma_0^2}{\sigma_0^2 + \sigma_u^2}$. Given the signal $y = \theta + u$ observed in the first period, $m_0 = \left(1 - K_0\right)\mu + K_0 y$. In period 2, the optimal continuation decision is

$$J(\theta) = \begin{cases} \frac{\theta}{1-\delta} & \text{if } \theta \geq \overline{\theta} \\ \delta W_0 & \text{if } \theta < \overline{\theta} \end{cases},$$

where $\overline{\theta} = (1-\delta)\,\delta W_0$ and W_0 denotes the expected present value for a worker starting with a new firm. Note that in period 2, conditional on m_0, we have that $\theta \sim N\left(m_0, \sigma_1^2\right)$, so that $\int_{-\infty}^{\infty} J\left(\theta'\right)\ \mathrm{d}G\left(\theta' | m_0, \sigma_1^2\right) = \int_{-\infty}^{\overline{\theta}} \delta W_0\ \mathrm{d}G\left(\theta' | m_0, \sigma_1^2\right) + \int_{\overline{\theta}(\mu)}^{\infty} \frac{\theta'}{1-\delta}\ \mathrm{d}G\left(\theta' | m_0, \sigma_1^2\right)$. As such

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}\mu} \int_{-\infty}^{\infty} J\left(\theta'\right)\ \mathrm{d}G\left(\theta' | m_0, \sigma_1^2\right) &= \delta \frac{\mathrm{d}W_0}{\mathrm{d}\mu} G\left(\overline{\theta} | m_0, \sigma_1^2\right) + \delta W_0 \frac{\mathrm{d}\overline{\theta}}{\mathrm{d}\mu} - \frac{\theta}{1-\delta} \frac{\mathrm{d}\overline{\theta}}{\mathrm{d}\mu} \\ &= \delta \frac{\mathrm{d}W_0}{\mathrm{d}\mu} G\left(\overline{\theta} | m_0, \sigma_1^2\right). \end{aligned}$$

We have that, conditional on y, $\theta \sim N\left(m_0, \sigma_1^2\right)$. The optimal continuation in period 1 is

$$V\left(m_0\right) = \begin{cases} m_0 + \delta \int_{-\infty}^{\infty} J\left(\theta'\right)\ \mathrm{d}G\left(\theta' | m_0, \sigma_1^2\right) & \text{if } m_0 \geq \overline{m_0} \\ \delta W_0 & \text{if } m_0 < \overline{m_0} \end{cases},$$

where $G\left(\cdot | m_0, \sigma_1^2\right)$ denotes the CDF of a normal distribution with mean m_0 and variance σ_1^2 and $\overline{m_0}$ solves

$$\delta W_0 = \overline{m_0} + \delta \int J\left(\theta'\right)\ \mathrm{d}G\left(\theta' | \overline{m_0}, \sigma_1^2\right). \tag{24}$$

We also have that

$$W_0 = \int V\left(m_0\right)\ \mathrm{d}G\left(m_0 | \mu, K_0 \sigma_0^2\right). \tag{25}$$

It is straightforward to see that a worker with a higher period 1 signal, and thus higher m_0, is less likely to quit in period 2, as he draws from a better θ distribution. Thus high type workers are less likely to quit in both periods. However, since our definition of task clarity is

[37] The two period version is sufficient to show that the comparative static goes in the opposite way.

ex ante and not interim, the most relevant comparison is across changes in μ. The effect here is less obvious—there is a direct effect of drawing higher productivities on average, but also an indirect effect of having a higher payoff from quitting the relationship. We show now that the direct effect dominates and that there is no "reversal" in the Jovanovic (1979) model.

Differentitating both sides of equation (24) with respect to μ we obtain

$$\begin{aligned}
\delta\frac{\mathrm{d}W_0}{\mathrm{d}\mu} &= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}+\delta\frac{\mathrm{d}}{\mathrm{d}\mu}\int_{-\infty}^{\infty}J\left(\theta'\right)\ \mathrm{d}G\left(\theta'|\overline{m_0},\sigma_1^2\right)\\
&= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}+\delta\frac{\mathrm{d}}{\mathrm{d}\mu}\int_{-\infty}^{\overline{\theta}(\mu)}\delta W_0 g\left(\theta'|\overline{m_0},\sigma_1^2\right)\ \mathrm{d}\theta'+\delta\frac{\mathrm{d}}{\mathrm{d}\mu}\int_{\overline{\theta}(\mu)}^{\infty}\frac{\theta'}{1-\delta}g\left(\theta'|\overline{m_0},\sigma_1^2\right)\ \mathrm{d}\theta'\\
&= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}+\delta^2\frac{\mathrm{d}W_0}{\mathrm{d}\mu}G\left(\overline{\theta}|\overline{m_0},\sigma_1^2\right)+\delta\int_{-\infty}^{\overline{\theta}(\mu)}\delta W_0\frac{\mathrm{d}}{\mathrm{d}\mu}g\left(\theta'|\overline{m_0},\sigma_1^2\right)\ \mathrm{d}\theta'\\
&\quad +\delta\int_{\overline{\theta}(\mu)}^{\infty}\frac{\theta'}{1-\delta}\frac{\mathrm{d}}{\mathrm{d}\mu}g\left(\theta'|\overline{m_0},\sigma_1^2\right)\ \mathrm{d}\theta'\\
&= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}+\delta^2\frac{\mathrm{d}W_0}{\mathrm{d}\mu}G\left(\overline{\theta}|\overline{m_0},\sigma_1^2\right)+\delta\frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}\int_{-\infty}^{\overline{\theta}(\mu)}\delta W_0\left(\frac{\theta'-\overline{m_0}}{\sigma_1^2}\right)g\left(\theta'|\overline{m_0},\sigma_1^2\right)\ \mathrm{d}\theta'\\
&\quad +\delta\frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}\int_{\overline{\theta}(\mu)}^{\infty}\frac{\theta'}{1-\delta}\left(\frac{\theta'-\overline{m_0}}{\sigma_1^2}\right)g\left(\theta'|\overline{m_0},\sigma_1^2\right)\ \mathrm{d}\theta'\\
&= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}+\delta^2\frac{\mathrm{d}W_0}{\mathrm{d}\mu}G\left(\overline{\theta}|\overline{m_0},\sigma_1^2\right)+\frac{\delta}{\sigma_1^2}\frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}\int_{-\infty}^{\infty}J\left(\theta'\right)\left(\theta'-\overline{m_0}\right)\ \mathrm{d}G\left(\theta'|\overline{m_0},\sigma_1^2\right)\\
&= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}+\delta^2\frac{\mathrm{d}W_0}{\mathrm{d}\mu}G\left(\overline{\theta}|\overline{m_0},\sigma_1^2\right)+\delta\frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}\int_{-\infty}^{\infty}\frac{\mathrm{d}}{\mathrm{d}\theta'}J\left(\theta'\right)\ \mathrm{d}G\left(\theta'|\overline{m_0},\sigma_1^2\right),
\end{aligned}$$

where the last line follows by Stein's lemma,[38] and we have canceled the boundary terms because of the definition of $\overline{\theta}(\mu)$. We can then write

$$\begin{aligned}
\frac{\mathrm{d}W_0}{\mathrm{d}\mu}\delta\left(1-\delta G\left(\overline{\theta}|\overline{m_0},\sigma_1^2\right)\right) &= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}\left(1+\delta\int_{-\infty}^{\infty}\frac{\mathrm{d}}{\mathrm{d}\theta'}J\left(\theta'\right)\ \mathrm{d}G\left(\theta'|\overline{m_0},\sigma_1^2\right)\right)\\
&= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}\left(1+\delta\int_{\overline{\theta}(\mu)}^{\infty}\frac{1}{1-\delta}\ \mathrm{d}G\left(\theta'|\overline{m_0},\sigma_1^2\right)\right)\\
&= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}\left(1+\frac{\delta}{1-\delta}\left(1-G\left(\overline{\theta}|\overline{m_0},\sigma_1^2\right)\right)\right)\\
&= \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}\left(\frac{1-\delta G\left(\overline{\theta}|\overline{m_0},\sigma_1^2\right)}{1-\delta}\right),
\end{aligned}$$

so that

$$\frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}=\frac{\mathrm{d}W_0}{\mathrm{d}\mu}\delta\left(1-\delta\right). \tag{26}$$

From equation (25) we have

$$\begin{aligned}
W_0 &= \int_{-\infty}^{\overline{m_0}}\delta W_0\ \mathrm{d}G\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)+\int_{\overline{m_0}}^{\infty}\left(m_0+\delta\int J\left(\theta'\right)\ \mathrm{d}G\left(\theta'|m_0,\sigma_1^2\right)\right)\ \mathrm{d}G\left(m_0|\mu,K_0\sigma_0^2\right)\\
&= \int_{-\infty}^{\overline{m_0}}\delta W_0 g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0+\int_{\overline{m_0}}^{\infty}\left(m_0+\delta\int J\left(\theta'\right)\ \mathrm{d}G\left(\theta'|m_0,\sigma_1^2\right)\right)g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0.
\end{aligned}$$

[38] Since J is only piecewise continuously differentiable, we need a generalization of the standard result, see Lemma 1.2 in Nourdin and Peccati (2009).

Differentiating both sides with respect to μ, noting that $\overline{m_0}$ depends on μ, yields

$$\begin{aligned}
\frac{\mathrm{d}W_0}{\mathrm{d}\mu} &= \delta\frac{\mathrm{d}W_0}{\mathrm{d}\mu}G\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)+\delta W_0 g\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)\frac{\mathrm{d}\overline{m_0}\left(\mu\right)}{\mathrm{d}\mu}+\int_{-\infty}^{\overline{m_0}}\delta W_0\frac{\partial}{\partial\mu}g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0 \\
&\quad -\left(\overline{m_0}+\delta\int J\left(\theta'\right)\ \mathrm{d}G\left(\theta'|\overline{m_0},\sigma_1^2\right)\right)g\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)\frac{\mathrm{d}\overline{m_0}\left(\mu\right)}{\mathrm{d}\mu} \\
&\quad +\int_{\overline{m_0}}^{\infty}\left(m_0+\delta\int J\left(\theta'\right)\ \mathrm{d}G\left(\theta'|m_0,\sigma_1^2\right)\right)\frac{\partial}{\partial\mu}g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0 \\
&\quad +\delta\int_{\overline{m_0}}^{\infty}\left(\delta\frac{\mathrm{d}W_0}{\mathrm{d}\mu}G\left(\overline{\theta}|m_0,\sigma_1^2\right)\right)g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0 \\
&= \delta\frac{\mathrm{d}W_0}{\mathrm{d}\mu}G\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)+\delta^2\frac{\mathrm{d}W_0}{\mathrm{d}\mu}\int_{\overline{m_0}}^{\infty}G\left(\overline{\theta}|m_0,\sigma_1^2\right)\ \mathrm{d}G\left(m_0|\mu,K_0\sigma_0^2\right) \\
&\quad +\int_{-\infty}^{\infty}V\left(m_0\right)\frac{\partial}{\partial\mu}g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0,
\end{aligned}$$

We can therefore write

$$\frac{\mathrm{d}W_0}{\mathrm{d}\mu}=\frac{\int_{-\infty}^{\infty}V\left(m_0\right)\frac{\partial}{\partial\mu}g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0}{1-\delta G\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)-\delta^2\int_{\overline{m_0}}^{\infty}G\left(\overline{\theta}|m_0,\sigma_1^2\right)g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0}. \tag{27}$$

Using Stein's lemma, we find that

$$\int_{-\infty}^{\infty}V\left(m_0\right)\frac{\partial}{\partial\mu}g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0 \quad=\quad \int_{-\infty}^{\infty}\frac{\mathrm{d}V\left(m_0\right)}{\mathrm{d}m_0}\ \mathrm{d}G\left(m_0|\mu,K_0\sigma_0^2\right).$$

Clearly $\frac{\mathrm{d}V(m_0)}{\mathrm{d}m_0}=0$ if $m_0<\overline{m_0}$. But for $m_0\geq\overline{m_0}$ we have

$$\begin{aligned}
\frac{\mathrm{d}V\left(m_0\right)}{\mathrm{d}m_0} &= 1+\delta\int_{-\infty}^{\infty}J\left(\theta'\right)\frac{\mathrm{d}}{\mathrm{d}m_0}g\left(\theta'|m_0,\sigma_1^2\right)\ \mathrm{d}\theta' \\
&= 1+\delta\int_{-\infty}^{\infty}J\left(\theta'\right)\frac{\theta'-m_0}{\sigma_1^2}g\left(\theta'|m_0,\sigma_1^2\right)\ \mathrm{d}\theta' \\
&= 1+\frac{\delta}{\sigma_1^2}\int_{-\infty}^{\infty}J\left(\theta'\right)\left(\theta'-m_0\right)\ \mathrm{d}G\left(\theta'|m_0,\sigma_1^2\right) \\
&= 1+\delta\int_{-\infty}^{\infty}\frac{\mathrm{d}}{\mathrm{d}\theta'}J\left(\theta'\right)\ \mathrm{d}G\left(\theta'|m_0,\sigma_1^2\right)=1+\delta\int_{\overline{\theta}}^{\infty}\frac{1}{1-\delta}\ \mathrm{d}G\left(\theta'|m_0,\sigma_1^2\right) \\
&= 1+\frac{\delta}{1-\delta}\left(1-G\left(\overline{\theta}|m_0,\sigma_1^2\right)\right)=\frac{1-\delta G\left(\overline{\theta}|m_0,\sigma_1^2\right)}{1-\delta},
\end{aligned}$$

where we again use Stein's lemma. Substituting back in to equation (27) gives

$$\begin{aligned}
\frac{\mathrm{d}W_0}{\mathrm{d}\mu} &= \frac{\int_{\overline{m_0}}^{\infty}\frac{1-\delta G\left(\overline{\theta}|m_0,\sigma_1^2\right)}{1-\delta}\ \mathrm{d}G\left(m_0|\mu,K_0\sigma_0^2\right)}{1-\delta G\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)-\delta^2\int_{\overline{m_0}}^{\infty}G\left(\overline{\theta}|m_0,\sigma_1^2\right)g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0} \\
&= \frac{1}{1-\delta}\frac{\int_{\overline{m_0}}^{\infty}\left(1-\delta G\left(\overline{\theta}|m_0,\sigma_1^2\right)\right)\ \mathrm{d}G\left(m_0|\mu,K_0\sigma_0^2\right)}{1-\delta G\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)-\delta^2\int_{\overline{m_0}}^{\infty}G\left(\overline{\theta}|m_0,\sigma_1^2\right)g\left(m_0|\mu,K_0\sigma_0^2\right)\ \mathrm{d}m_0} \\
&= \frac{1}{1-\delta}\frac{1-G\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)-\delta\int_{\overline{m_0}}^{\infty}G\left(\overline{\theta}|m_0,\sigma_1^2\right)\ \mathrm{d}G\left(m_0|\mu,K_0\sigma_0^2\right)}{1-\delta G\left(\overline{m_0}|\mu,K_0\sigma_0^2\right)-\delta^2\int_{\overline{m_0}}^{\infty}G\left(\overline{\theta}|m_0,\sigma_1^2\right)\ \mathrm{d}G\left(m_0|\mu,K_0\sigma_0^2\right)}.
\end{aligned}$$

To see that $\frac{\mathrm{d}W_0}{\mathrm{d}\mu} \leq \frac{1}{1-\delta}$ note that the denominator of the second factor is subtracting a number δ times smaller than the numerator. Also, both the numerator and denominator are positive because $G\left(\overline{m_0}|\mu, K_0\sigma_0^2\right) + \delta \int_{\overline{m_0}}^{\infty} G\left(\overline{\theta}|m_0, \sigma_1^2\right)\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) < G\left(\overline{m_0}|\mu, K_0\sigma_0^2\right) + \int_{\overline{m_0}}^{\infty} 1\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) = 1$. Hence $\frac{\mathrm{d}W_0}{\mathrm{d}\mu} > 0$.

Finally, we have that the probability that an unemployed worker accepts an offer is $\mathbb{P}\left[m_0 \geq \overline{m_0}\right] = 1 - G\left(\overline{m_0}|\mu, K_0\sigma_0^2\right)$. Differentiating this with respect to μ, and using equation (26) we get

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}\mu}\mathbb{P}\left[m_0 \geq \overline{m_0}\right] &= -g\left(\overline{m_0}\left(\mu\right)|\mu, K_0\sigma_0^2\right)\frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu} + g\left(\overline{m_0}\left(\mu\right)|\mu, K_0\sigma_0^2\right) \\
&= g\left(\overline{m_0}\left(\mu\right)|\mu, K_0\sigma_0^2\right)\left(1 - \frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu}\right) \\
&= g\left(\overline{m_0}\left(\mu\right)|\mu, K_0\sigma_0^2\right)\left(1 - \frac{\mathrm{d}W_0}{\mathrm{d}\mu}\delta\left(1-\delta\right)\right) \\
&\geq g\left(\overline{m_0}\left(\mu\right)|\mu, K_0\sigma_0^2\right)\left(1-\delta\right) > 0.
\end{aligned}$$

So, workers with a higher mean match productivity are more likely to accept an offer in the first period. This is despite the fact that the cutoff for accepting offers is also higher.

We are also interested in the probability that a worker quits the relationship in period 2, conditional on having stayed in period 1. This is

$$\frac{\mathbb{P}\left[\theta < \overline{\theta}, m_0 \geq \overline{m_0}\right]}{\mathbb{P}\left[m_0 \geq \overline{m_0}\right]} = \frac{\int_{\overline{m_0}}^{\infty} G\left(\overline{\theta}|m_0, \sigma_1^2\right) g\left(m_0|\mu, K_0\sigma_0^2\right)\ \mathrm{d}m_0}{1 - G\left(\overline{m_0}|\mu, K_0\sigma_0^2\right)}.$$

We first consider the derivative of the numerator with respect to μ, namely the joint probability that a worker quits in period 2, and accepts in period 1, and show that it is negative. We have

$$\begin{aligned}
&\frac{\mathrm{d}}{\mathrm{d}\mu}\mathbb{P}\left[\theta < \overline{\theta}, m_0 \geq \overline{m_0}\right] = \frac{\mathrm{d}}{\mathrm{d}\mu}\int_{\overline{m_0}}^{\infty} G\left(\overline{\theta}|m_0, \sigma_1^2\right)\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) \\
= \;& -G\left(\overline{\theta}|\overline{m_0}, \sigma_1^2\right) g\left(\overline{m_0}|\mu, K_0\sigma_0^2\right)\frac{\mathrm{d}\overline{m_0}}{\mathrm{d}\mu} \\
&+ \int_{\overline{m_0}}^{\infty} g\left(\overline{\theta}|m_0, \sigma_1^2\right)\frac{\mathrm{d}\overline{\theta}}{\mathrm{d}\mu}\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) + \int_{\overline{m_0}}^{\infty} G\left(\overline{\theta}|m_0, \sigma_1^2\right)\frac{\mathrm{d}}{\mathrm{d}\mu}g\left(m_0|\mu, K_0\sigma_0^2\right)\ \mathrm{d}m_0 \\
= \;& -\delta\left(1-\delta\right) G\left(\overline{\theta}|\overline{m_0}, \sigma_1^2\right) g\left(\overline{m_0}|\mu, K_0\sigma_0^2\right)\frac{\mathrm{d}W_0}{\mathrm{d}\mu} \\
&+ \int_{\overline{m_0}}^{\infty} g\left(\overline{\theta}|m_0, \sigma_1^2\right)\delta\left(1-\delta\right)\frac{\mathrm{d}W_0}{\mathrm{d}\mu}\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) + \int_{\overline{m_0}}^{\infty} \frac{\mathrm{d}}{\mathrm{d}m_0}G\left(\overline{\theta}|m_0, \sigma_1^2\right)\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) \\
= \;& -\delta\left(1-\delta\right) G\left(\overline{\theta}|\overline{m_0}, \sigma_1^2\right) g\left(\overline{m_0}|\mu, K_0\sigma_0^2\right)\frac{\mathrm{d}W_0}{\mathrm{d}\mu} \\
&+ \int_{\overline{m_0}}^{\infty} g\left(\overline{\theta}|m_0, \sigma_1^2\right)\delta\left(1-\delta\right)\frac{\mathrm{d}W_0}{\mathrm{d}\mu}\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) - \int_{\overline{m_0}}^{\infty} g\left(\overline{\theta}|m_0, \sigma_1^2\right)\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) \\
= \;& -\delta\left(1-\delta\right) G\left(\overline{\theta}|\overline{m_0}, \sigma_1^2\right) g\left(\overline{m_0}|\mu, K_0\sigma_0^2\right)\frac{\mathrm{d}W_0}{\mathrm{d}\mu} \\
&+ \int_{\overline{m_0}}^{\infty} \left(\delta\left(1-\delta\right)\frac{\mathrm{d}W_0}{\mathrm{d}\mu} - 1\right) g\left(\overline{\theta}|m_0, \sigma_1^2\right)\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) \\
\leq \;& -\delta\left(1-\delta\right) G\left(\overline{\theta}|\overline{m_0}, \sigma_1^2\right) g\left(\overline{m_0}|\mu, K_0\sigma_0^2\right)\frac{\mathrm{d}W_0}{\mathrm{d}\mu} + \int_{\overline{m_0}}^{\infty} \left(\delta - 1\right) g\left(\overline{\theta}|m_0, \sigma_1^2\right)\ \mathrm{d}G\left(m_0|\mu, K_0\sigma_0^2\right) \\
< \;& 0.
\end{aligned}$$

In order to find the sign on how the conditional probability of quitting in period 2 changes with respect to μ, given acceptance in period 1, $\frac{\mathbb{P}\left[\theta<\overline{\theta}, m_0 \geq \overline{m_0}\right]}{\mathbb{P}\left[m_0 \geq \overline{m_0}\right]}$, we apply the quotient rule. The sign of this is given by the sign of

$$\mathbb{P}\left[m_0 \geq \overline{m_0}\right] \frac{\mathrm{d}}{\mathrm{d}\mu} \mathbb{P}\left[\theta < \overline{\theta}, m_0 \geq \overline{m_0}\right] - \mathbb{P}\left[\theta < \overline{\theta}, m_0 \geq \overline{m_0}\right] \frac{\mathrm{d}}{\mathrm{d}\mu} \mathbb{P}\left[m_0 \geq \overline{m_0}\right],$$

but note that both terms are negative. The first is negative since $\frac{\mathrm{d}}{\mathrm{d}\mu} \mathbb{P}\left[\theta < \overline{\theta}, m_0 \geq \overline{m_0}\right] < 0$ as we have just shown and the second is negative since we previously established that $\frac{\mathrm{d}}{\mathrm{d}\mu} \mathbb{P}\left[m_0 \geq \overline{m_0}\right] > 0$.

Thus, as μ increases the conditional probability of quitting in period 2 is lower, and hence the conditional probability of staying in the relationship in period 2 is higher.

C Additional Tables

Table C.1: Relational Statistics

Dependent Variable:	All Shipments (1)	Direct Shipments (2)	Auction Shipments (3)	Sell Direct (%) (4)	Relationships Attempted (5)	Success (%) (6)	Relationships Ended (7)
I[Domestic Exp.]	-12.423	-99.578***	87.155***	-0.465***	-16.334*	-0.157*	-18.652*
	(40.982)	(29.172)	(24.251)	(0.087)	(8.598)	(0.084)	(9.664)
Mean Foreign	183.345	137.916	45.429	0.657	34.818	0.871	37.523
Observations	75	75	75	74	75	64	75

Note: The table compares various relational statistics between foreign and domestically owned exporters using customs data. Estimates are based on the following firm-level regression: $Outcome_f = \beta_0 + \beta_1 I[Domestic]_f + \epsilon_f$, where $Outcome_f$ is a firm-level outcome variable from the data, and $I[Domestic]_f$ is an indicator variable equal to 1 if firm f is a domestically owned exporter, and 0 otherwise. Outcomes include several key variables used in the empirical analysis. *All Shipments* is the average number of shipments per season at the exporter level. This measure is further disaggregated by export channel: *Direct* refers to transactions made through relationships, while *Auction* refers to transactions sold via the auction house. *Share Sell Direct* is the average proportion of each exporter's sales conducted through relationships. *Relationships Attempted* is the mean number of buyer relationships initiated by an exporter (i.e., where at least one shipment was sent). *Share Success* is the average rate at which attempted relationships reached at least a fourth shipment. *Relationships Ended* is the mean number of direct relationships that terminated, defined as no shipments between a buyer and a seller for at least six months. The sample includes 44 foreign-owned and 31 domestically owned exporters. Robust standard errors are reported in parentheses. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 levels, respectively.

Table C.2: Differences across Shipments

Dependant Variable:	Shipment Size:					
	Quantity (no. of stems)	Weight (kgs)	Value (USD)	Quantity (no. of stems)	Weight (kgs)	Value (USD)
	(1)	(2)	(3)	(4)	(5)	(6)
I[First 3 Shipments]	-767.683*	-33.854***	-136.352***	-770.657**	-33.943***	-136.021***
	(385.009)	(10.693)	(41.430)	(384.359)	(10.689)	(41.383)
Mean Dep. Var	20,440.133	712.657	2,660.560	20,265.074	707.414	2,641.174
% Relative to Mean	-3.8%	-4.8%	-5.1%	-3.8%	-4.8%	-5.2%
Observations	64,872	64,872	63,664	65,765	65,765	64,479
Controls:						
Month x Year FE	Y	Y	Y	Y	Y	Y
Exporter-Buyer Pair FE	Y	Y	Y	Y	Y	Y
Productive Relationships	Y	Y	Y			

Note: The table estimates how shipment characteristics – quantity (stems), weight (kg) and value (USD $) – between between the first three shipments in a buyer-seller relationship and all subsequent shipments, using: $Y_{s,p,t} = \psi_p + \zeta_t + \beta \mathbb{1}[s \leq 3]_{s,p} + \epsilon_{s,p,t}$. Where $Y_{s,p,t}$ is the outcome for shipment s in pair p at time t, ψ_p denotes buyer-seller pair (relationship) fixed effects, ζ_t are month-by-year fixed effects, and $\mathbb{1}[s \leq 3]_{s,p}$ in an indicator variable equal to 1 for the first three shipments in a relationship, and 0 otherwise. Columns 1-3 restrict to productive relationships (those that eventually reach a fourth shipment), while columns 4-6 include all relationships. Standard errors, in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.3: Varieties of Roses Cultivated

Dependent Variable:	Land (hectares)			Roses (% of Total Land)			I[Produces Variety]			Average Yield (stems/m^2)		
	SW	IM	T-H	SW	IM	T-H	SW	IM	T-H	SW	IM	T-H
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
I[Domestic Exp.]	1.676	-0.073	-0.432	0.088	-0.018	-0.061	0.056	0.314**	0.126	44.500	6.292	2.778
	(1.266)	(1.476)	(0.757)	(0.082)	(0.115)	(0.117)	(0.127)	(0.122)	(0.153)	(45.657)	(15.208)	(13.559)
Mean Dep. Var	0.883	4.307	2.245	0.073	0.546	0.383	0.173	0.642	0.509	230.333	167.879	129.259
Observations	52	53	53	40	41	41	52	53	53	9	33	27

Note: The table compares rose variety production between foreign- and domestic-owned exporters using the 2008 firm census. The estimates are based on the following firm-level regression: $Outcome_f = \beta_0 + \beta_1 I[Domestic]_f + \epsilon_f$, where $Outcome_f$ is a firm-level outcome variable from the census, and $I[Domestic]_f$ is an indicator variable equal to 1 if firm f is a domestically owned exporter, 0 otherwise. A coarse proxy for quality is the rose variety, which is ranked from smallest and least expensive to largest and most expensive: Sweetheart, Intermediate, and Tea-hybrid. The census reports these varieties at the firm-level. Columns 1-3 report differences in land allocation (in hectares) across the three rose varieties: Sweetheart (SW), Intermediate (IM), and Tea-Hybrid (T-H). Columns 4-6 display the share of total rose cultivation land allocated to each variety. Columns 7-9 present dummy indicators for whether the firm produces each variety. Columns 10-12 report the average yield (stems per square meter) for each variety, conditional on the firm producing that variety. The unit of observation is the number of firms surveyed in the 2008 census. Robust standard errors are reported in parentheses. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.4: Rose Quality

Dependant Variable:	Unit stem weight (ln, kg)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
I[Domestic Exp.]	0.032	0.035	0.037	0.108	0.092	0.089	-0.038	-0.003	0.000
	(0.105)	(0.098)	(0.096)	(0.149)	(0.142)	(0.134)	(0.089)	(0.082)	(0.082)
Sample	All Shipments			Direct Relationships			Auction Sales		
Observations	120,937	120,937	120,937	66,517	66,517	66,517	54,420	54,420	54,420
Season FE		Y			Y			Y	
Month x Year FE			Y			Y			Y

Note: The table compares the rose quality measured, by unit stem weight (kg) between foreign- and domestic-owned exporters using the transaction data. The estimates are based on the following shipment-level regression: $Log(UnitStemWeight)_{s,t} = \beta I[Domestic]_s + \zeta_t + \epsilon_{s,t}$, where $Log(UnitStemWeight)_{s,t}$ is the stem unit weight (natural logarithm, kg) of shipment s during week t, and $I[Domestic]_f$ is an indicator variable equal to 1 if the shipment is from a domestically owned exporter, 0 otherwise. Shipments are weighted by the number of flower stems they contain, so the comparison reflects the average flower weight (kg). Season Fixed Effects and month-by-year fixed effects are included as captured by ζ_t. In Columns 1-3, the dependent variable includes all shipments. Columns 4–6 restrict the sample to direct shipments only, while Columns 7–9 are limited to auction-only shipments. Standard errors, in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.5: Production Process for Cultivation of Roses

Dependent Variable:	Altitude (meters)	Land (hectares) under Greenhouses	Rejection Rate (%)	Importance for quality control (scale 0-4)					% Local	
				Feedback	Seedlings	Phyto. Insp.	Sorting	Ext. Insp.	Fertilizers	Chemicals
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
I[Domestic Exp.]	82.833	-1.100	-0.321	0.551***	-0.181	0.018	0.447***	0.319	19.577	16.063
	(128.668)	(1.795)	(0.907)	(0.200)	(0.392)	(0.338)	(0.135)	(0.452)	(14.507)	(14.777)
Mean Dep. Var	2,064.943	8.525	2.893	3.472	3.396	3.321	3.679	2.038	32.404	34.904
Observations	53	53	51	53	53	53	53	53	52	52

Note: The table compares various farm characteristics and quality control practices between foreign and domestically owned exporters using the 2008 firm census. The estimates are based on the following firm-level regression: $Outcome_f = \beta_0 + \beta_1 I[Domestic]_f + \epsilon_f$, where $Outcome_f$ is a firm-level outcome variable from the census, and $I[Domestic]_f$ is an indicator variable equal to 1 if firm f is a domestically owned exporter, 0 otherwise. Column 1 indicates the altitude of the farm (in meters above sea level). Column 2 shows the land covered by greenhouses (in hectares) for flower production. Column 3 compares the rejection rate (% in terms of quantity) at export destination. Rejection rate was winsorized at 5 and 95 percentiles. Columns 4-8 assess the importance of various quality control measures in the flower production process, rated on a scale from 0 (not important) to 4 (critical importance). Column 4 evaluates the presence of quality control staff who gather feedback on flower quality from buyers. Column 5 focuses on the use of high-quality seedlings. Column 6 addresses the conduct of phytosanitary inspections to ensure pre- and post-harvest quality control. Column 7 looks at the regular sorting of damaged and diseased flowers, while Column 8 examines the practice of inviting external audits for periodic inspections. Columns 9 and 10 report the percentage of local manufactured fertilizers and chemicals used in 2007, respectively. The unit of observation is the number of firms surveyed in the 2008 census. Robust standard errors are reported in parentheses. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.6: Cost of Capital and Discount Factor

Dependent Variable:	External Funds (% of working capital)		Hardship with Credit (Yes=1)		Share of Collateral (% of Total Loan)		Interest Rate (%)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
I[Domestic Exp.]	0.204**	0.256**	0.340**	0.583***	0.451**	0.498*	0.858	1.066*
	(0.089)	(0.101)	(0.150)	(0.154)	(0.218)	(0.260)	(0.555)	(0.588)
Mean Foreign	0.325	0.309	0.393	0.359	0.521	0.540	8.356	8.252
Observations	48	46	43	42	69	46	39	38
Controls		Y		Y		Y		Y

Note: The table presents the difference in reliance on credit, access to credit, and its cost between foreign- and domestic-owned exporters, based on data from the 2008 and 2011 firm censuses. The estimates are based on the following firm-level OLS regression: $Outcome_f = \beta_1 I[Domestic]_f + \Gamma X_f + \epsilon_f$, where $Outcome_f$ is a firm-level outcome variable from the census, and $I[Domestic]_f$ is an indicator variable equal to 1 if firm f is a domestically owned exporter, 0 otherwise. Differences in sample size across columns arise due to missing data for either the outcome variables or the control variables. In Columns 1–2, the dependent variable is External funds defined as share of working capital not derived from internal funds or retained earnings in 2008. In Columns 3–4, the dependent variable is Hardship with credit, defined as a binary variable equal to 1 if the firm indicates that both access to credit and the cost of financing (e.g., interest rates) constitute a moderate, major or very severe obstacle to business operations or growth. In Columns 5–6, the dependent variable is the Share of collateral, defined as the value of collateral required, expressed as a percentage of the loan value. In Columns 7–8, the dependent variable is the firm's marginal interest rate, defined as the highest rate paid on short- or long-term liabilities. Columns 2, 4, 6, and 8 include control variables. These controls consist of the first three principal components (PCs) derived from a PCA conducted on the following 2008 survey variables: total owned land (ha), land under flower cultivation (ha), land under greenhouse production (ha), land covered under roses (ha), distance from the farm to the airport (km), weekly number of flower stems delivered to the airport, total number of workers, number of full-time employees, number of production workers in 2007, and joint venture status. The variables for external funds, hardship with credit, and interest rate are sourced from the 2008 firm survey, the corresponding questions were not asked in the 2011 census. The share of collateral is computed as the average of the firm's responses from both 2008 and 2011. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.7: Domestically Owned Exporters and Ending Productive Relationships

Dependent Variable:	Number of Relationships Ending ≥ 1 (Dummy)		Number of Relationships Ending (Count)			
	(1)	(2)	(3)	(4)	(5)	(6)
Price Spread (Std)	0.010**	0.010**	0.012**	0.012**	0.138*	0.137*
	(0.004)	(0.004)	(0.005)	(0.005)	(0.076)	(0.076)
Price Spread (Std) x	-0.010*	-0.009*	-0.010*	-0.010	-0.328*	-0.325*
I[Domestic Exp.]	(0.005)	(0.005)	(0.006)	(0.006)	(0.176)	(0.176)
Mean Dep. Var	0.047	0.047	0.053	0.053	0.059	0.059
Observations	17,488	17,488	17,488	17,488	15,611	15,611
Controls:						
Productive Rel. (N)	Y	Y	Y	Y	Y	Y
Quantity (Ln, no. of stems)	Y	Y	Y	Y	Y	Y
Value (Ln, USD)		Y		Y		Y
Exporter FE	Y	Y	Y	Y	Y	Y
Estimation	OLS	OLS	OLS	OLS	PPML	PPML

Note: The table reports OLS and PPML estimates of β_2 and β_3 at the exporter-week level: $Y_{e,t} = \beta_0 + \beta_1 D_e + \beta_2 Price\ Spread_{t+1} + \beta_3 Price\ Spread_{t+1} \times D_e + \Gamma X_{e,t} + \epsilon_{e,t}$. Where the dependent variable, $Y_{e,t}$, in Columns 1-2, is a dummy equal to one if the exporter experienced at least one relationship termination in the week, while in Columns 3-6, it is the count of relationships terminating that week. Columns 5-6 are estimated using a Poisson Pseudo-Maximum Likelihood (PPML) estimator. A relationship is deemed to end if no further shipments occur between the exporter and a given seller for six months. *Price Spread* is the standardized difference between the average auction price and the average price direct-shipment price in the subsequent week. The variable *Domestic Exp.* is a dummy indicating whether the exporter is domestically owned or not. The sample includes only *productive* relationships (those that survived beyond the third shipment). Controls (which vary across specifications) include exporter fixed effects, current number of active productive relationships (Productive Rel.), weekly shipment quantity and the value of the exporters' transactions (in logs). Standard errors, in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at a 0.10, 0.05, and 0.01 level, respectively.

Table C.8: Productive Relationships and Exporter- and Buyer-Specific Components

Productive:	Reach 3 Shipments	Reach 4 Shipments	Reach 5 Shipments	Reach 6 Shipments
	(1)	(2)	(3)	(4)
Exporter Component (θ_e)	0.216***	0.214***	0.231***	0.229***
	(0.019)	(0.016)	(0.017)	(0.014)
	[19.05%]	[16.32%]	[18.61%]	[20.13%]
Buyer Component (θ_b)	0.277***	0.290***	0.279***	0.275***
	(0.017)	(0.017)	(0.017)	(0.014)
	[27.20%]	[31.94%]	[29.14%]	[28.76%]
Mean Dep. Var	0.527	0.458	0.411	0.385
Observations	1,378	1,378	1,378	1,378

Note: The table reports OLS estimates from the regression: $\mathbb{1}[Productive]_{e,b} = \beta_0 + \beta_1\hat{\theta_e} + \beta_2\hat{\theta_b} + \epsilon_{e,b}$. Differences across columns reflect alternative definitions of a productive relationship – that is, the number of shipments required to classify a relationship as productive. Exporter and buyer components are estimated using the OLS regression: $\mathbb{1}[Productive]_{e,b} = \theta_e + \theta_b + \varepsilon_{e,b}$, then standardized and winsorized at the 5% level. Standard errors, in parentheses, are clustered at the exporter level. Share of the variance of the outcome explained by the variable, in brackets, estimated using an Analysis of Variance (ANOVA). Asterisks next to the estimate, *, **, ***, denote statistical significance at a 0.10, 0.05, and 0.01 level, respectively. Observations are weighted so that each exporter receives equal weight.

Table C.9: Exporter Component and Ending Productive Relationships, Controls

Dependent Variable:	Number of Relationships Ending ≥ 1 (Dummy)		Number of Relationships Ending (Count)			
	(1)	(2)	(3)	(4)	(5)	(6)
Price Spread (Std)	0.010***	0.010***	0.012***	0.012***	0.024	0.024
	(0.003)	(0.003)	(0.004)	(0.004)	(0.067)	(0.068)
Price Spread (Std) x	0.012***	0.012***	0.013***	0.013***	0.221**	0.220**
Exporter Comp. (θ_e)	(0.003)	(0.003)	(0.004)	(0.004)	(0.087)	(0.087)
Mean Dep. Var	0.052	0.052	0.059	0.059	0.060	0.060
Observations	15,585	15,585	15,585	15,585	15,329	15,329
Controls:						
Productive Rel. (N)	Y	Y	Y	Y	Y	Y
Quantity (Ln, no. of stems)	Y	Y	Y	Y	Y	Y
Value (Ln, USD)		Y		Y		Y
Exporter FE	Y	Y	Y	Y	Y	Y
Estimation	OLS	OLS	OLS	OLS	PPML	PPML

Note: The table reports OLS and PPML estimates of β_2 and β_3 from the following exporter-week level regression: $Y_{e,t} = \beta_0 + \beta_1\hat{\theta_e} + \beta_2 Price\ Spread_{t+1} + \beta_3 Price\ Spread_{t+1} \times \hat{\theta_e} + \Gamma X_{e,t} + \epsilon_{e,t}$. In Columns 1-2, the dependent variable is a dummy equal to one if the exporter had at least one relationship end in a given week. In Columns 3-6, the dependent variable is the count of relationships that ended. Columns 5-6 are estimated using a Poisson Pseudo-Maximum Likelihood (PPML) estimator. A relationship is considered ended if no shipments occur between an exporter and a buyer for at least six months. *Price Spread* is defined as the standardized difference between the average auction price and the average price in direct shipments for the following week. The exporter component, $\hat{\theta_e}$, is estimated using regression: $\mathbb{1}[Productive]_{e,b} = \theta_e + \theta_b + \varepsilon_{e,b}$, standardized, and winsorized at the 5% level. The sample includes all productive relationships (those that survived beyond the third shipment). Control variables vary across specifications and include exporter fixed effects, the number of active productive relationships, and the weekly quantity and value of the exporters' weekly transactions (both in logs). Standard errors, in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.10: Robustness for Outliers in Price Distribution

Dependent Variable:	Number of Relationships Ending ≥ 1 (Dummy)			Number of Relationships Ending (Count)					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Price Spread (Std)	0.011***	0.011***	0.011***	0.012***	0.013***	0.013***	0.024	0.026	0.028
	(0.003)	(0.004)	(0.003)	(0.004)	(0.005)	(0.004)	(0.067)	(0.071)	(0.068)
Price Spread (Std) x	0.012***	0.012***	0.012***	0.013***	0.014***	0.013***	0.221**	0.229**	0.219**
Exporter Comp. (θ_e)	(0.003)	(0.004)	(0.003)	(0.004)	(0.005)	(0.004)	(0.086)	(0.093)	(0.088)
Mean Dep. Var	0.052	0.052	0.052	0.059	0.059	0.059	0.060	0.060	0.060
Observations	15,585	15,585	15,585	15,585	15,585	15,585	15,329	15,329	15,329
Sample	All	Trim [5-95]	Winsor [5-95]	All	Trim [5-95]	Winsor [5-95]	All	Trim [5-95]	Winsor [5-95]
Estimation	OLS	OLS	OLS	OLS	OLS	OLS	PPML	PPML	PPML

Note: The table reports OLS and PPML estimates of β_2 and β_3 at the exporter-week level from equation $Y_{e,t} = \beta_0 + \beta_1 \hat{\theta_e} + \beta_2 Price\ Spread_{t+1} + \beta_3 Price\ Spread_{t+1} \times \hat{\theta_e} + \Gamma X_{e,t} + \epsilon_{e,t}$. In Columns 1-3, the dependent variable is a binary indicator equal to one if the exporter had at least one relationship ending in that week. In Columns 4-9, the dependent variable is the number of relationships that ended in the week, with Columns 7-9 estimated using a Poisson Pseudo-Maximum Likelihood (PPML) estimator. A relationship is defined as ending if no shipments are observed between a given exporter and seller for at least six months. *Price Spread* is measured as the standardized difference between the average auction price and the average direct-shipment price for the following week. The exporter component, $\hat{\theta_e}$, is estimated using equation $\mathbb{1}[Productive]_{e,b} = \theta_e + \theta_b + \varepsilon_{e,b}$, standardized, and winsorized at the 5% level. The sample includes all *productive* relationships – those surviving past the third shipment. All specifications control for the number of active, productive relationships the exporter has in a given week and include exporter fixed effects. Differences across columns arise from winsorizing (Columns 3, 6, and 9) or trimming (Columns 2, 5, and 8) of the shipment's unit price (USD). Standard errors, in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.11: Exporter Component and Ending Productive Relationships (9 months)

Dependent Variable:	Number of Relationships Ending ≥ 1 (Dummy)		Number of Relationships Ending (Count)			
	(1)	(2)	(3)	(4)	(5)	(6)
Price Spread (Std)	0.007**	0.008***	0.008**	0.009**	0.018	-0.005
	(0.003)	(0.003)	(0.004)	(0.003)	(0.067)	(0.066)
Price Spread (Std) x	0.007**	0.009***	0.007**	0.010***	0.186***	0.202**
Exporter Comp. (θ_e)	(0.003)	(0.002)	(0.003)	(0.003)	(0.065)	(0.081)
Mean Dep. Var	0.050	0.050	0.055	0.055	0.055	0.056
Observations	15,248	15,248	15,248	15,248	15,248	15,006
Exporter FE		Y		Y		Y
Estimation	OLS	OLS	OLS	OLS	PPML	PPML

Note: The table reports the OLS and PPML estimates of β_2 and β_3 at the exporter-week level from equation $Y_{e,t} = \beta_0 + \beta_1 \hat{\theta_e} + \beta_2 Price\ Spread_{t+1} + \beta_3 Price\ Spread_{t+1} \times \hat{\theta_e} + \Gamma X_{e,t} + \epsilon_{e,t}$. In Columns 1-2, the dependent variable is a binary indicator that equals one if the exporter had at least one relationship ending in that week. In Columns 3-6, the dependent variable is the number of relationships that ended in the week, with Columns 5-6 estimated using a Poisson Pseudo-Maximum Likelihood (PPML) estimator. A relationship is considered to have ended if no shipments are observed between an exporter and a seller for at least *nine months*. *Price Spread* is defined as the standardized difference between the average auction price and the average direct-shipment price for the following week. The exporter component, $\hat{\theta_e}$, is estimated using equation $\lambda(\theta_e, \theta_b) = \frac{1}{1+e^{-(\theta_e+\theta_b)}}$, standardized, and winsorized at the 5% level. The sample includes all productive relationships (survived past the third shipment). All specifications include a control for the number of active, productive relationships the exporter has that week. Standard errors, in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.12: Exporter Component and Ending Productive Relationships, Post-2012

Dependent Variable:	Number of Relationships Ending ≥ 1 (Dummy)		Number of Relationships Ending (Count)			
	(1)	(2)	(3)	(4)	(5)	(6)
Price Spread (Std)	0.010	0.011	0.013	0.014	0.008	-0.027
	(0.007)	(0.007)	(0.009)	(0.010)	(0.085)	(0.076)
Price Spread (Std) x	0.013*	0.017*	0.015	0.019*	0.253**	0.294**
Exporter Comp. (θ_e)	(0.007)	(0.009)	(0.009)	(0.011)	(0.102)	(0.115)
Mean Dep. Var	0.061	0.061	0.070	0.070	0.070	0.074
Observations	10,276	10,276	10,276	10,276	10,276	9,641
Exporter FE		Y		Y		Y
Estimation	OLS	OLS	OLS	OLS	PPML	PPML

Note: The table reports the OLS and PPML estimates of β_2 and β_3 at the exporter-week level from equation $Y_{e,t} = \beta_0 + \beta_1 \hat{\theta}_e + \beta_2 Price\ Spread_{t+1} + \beta_3 Price\ Spread_{t+1} \times \hat{\theta}_e + \Gamma X_{e,t} + \epsilon_{e,t}$, using only shipments from 2012 onward. In Columns 1-2, the dependent variable is a binary indicator equal to one if the exporter had at least one relationship ending in that week. In Columns 3-6, the dependent variable is number of relationships that ended in the week, with Columns 5-6 estimated using a Poisson Pseudo-Maximum Likelihood (PPML) estimator. A relationship is considered to have ended if no shipments are observed between an exporter and a seller for at least six months. *Price Spread* is defined as the standardized difference between the average auction price and the average direct-shipment price for the following week. The exporter component, $\hat{\theta}_e$, is estimated using equation $\mathbb{1}[Productive]_{e,b} = \theta_e + \theta_b + \varepsilon_{e,b}$, standardized, and winsorized at the 5% level. The sample includes all *productive* relationships – those that survived beyond the third shipment. All specifications control for the number of active, productive relationships the exporter has in a given week. Standard errors, in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.13: Exporter Component and Ending Productive Relationships, Monthly Frequency

Dependent Variable:	Number of Relationships Ending ≥ 1 (Dummy)		Number of Relationships Ending (Count)			
	(1)	(2)	(3)	(4)	(5)	(6)
Price Spread (Std)	0.021**	0.021**	0.032**	0.036**	0.036	0.014
	(0.009)	(0.008)	(0.013)	(0.014)	(0.056)	(0.060)
Price Spread (Std) x	0.021*	0.031***	0.027**	0.045***	0.148**	0.148**
Exporter Comp. (θ_e)	(0.011)	(0.011)	(0.012)	(0.014)	(0.061)	(0.071)
Mean Dep. Var	0.220	0.220	0.309	0.309	0.309	0.313
Observations	2,965	2,965	2,965	2,965	2,965	2,927
Exporter FE		Y		Y		Y
Estimation	OLS	OLS	OLS	OLS	PPML	PPML

Note: The table reports the OLS and PPML estimates of β_2 and β_3 at the exporter-month from equation $Y_{e,t} = \beta_0 + \beta_1 \hat{\theta}_e + \beta_2 Price\ Spread_{t+1} + \beta_3 Price\ Spread_{t+1} \times \hat{\theta}_e + \Gamma X_{e,t} + \epsilon_{e,t}$. In Columns 1-2, the dependent variable is a binary indicator equal to one if the exporter had at least one relationship ending in that month. In Columns 3-6, the dependent variable is the number of relationships that ended in the month, with Columns 5-6 estimated using a Poisson Pseudo-Maximum Likelihood (PPML) estimator. A relationship is defined as ending if no shipments are observed between an exporter and a seller for at least six months. *Price Spread* is measured as the standardized difference between the average auction price and the average direct-shipment price for that month. The exporter component, $\hat{\theta}_e$, is estimated using equation $\mathbb{1}[Productive]_{e,b} = \theta_e + \theta_b + \varepsilon_{e,b}$, standardized, and winsorized at the 5% level. The sample includes all *productive* relationships – those surviving beyond the third shipment. All specifications control for the number of active, productive relationships the exporter has in that month. Standard errors, in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

Table C.14: Domestic Firms and Attempting New Relationships

Dependent Variable:	Number of Attempts ≥ 1 (Dummy)			Number of Attempts (Count)					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
I[Domestic Exp.]	-0.050**	-0.054**	-0.053**	-0.067**	-0.074**	-0.073**	-0.640***	-0.692***	-0.685***
	(0.020)	(0.024)	(0.024)	(0.028)	(0.033)	(0.033)	(0.234)	(0.260)	(0.256)
Mean Dep. Var	0.095	0.095	0.095	0.114	0.114	0.114	0.129	0.129	0.129
Observations	18,428	18,428	18,428	18,428	18,428	18,428	16,362	16,362	16,362
Controls:									
Quantity (Ln, no. of stems)		Y	Y		Y	Y		Y	Y
Value (Ln, USD)			Y			Y			Y
Week x Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Estimation	OLS	OLS	OLS	OLS	OLS	OLS	PPML	PPML	PPML

Note: The table reports estimates at the exporter-week level from the equation $Y_{e,t} = \beta_0 + \beta_1 Domestic + \Gamma X_{e,t} + \zeta_t + \epsilon_{e,t}$. For Columns 1-3, the dependent variable is a binary indicator equal to one if the exporter made at least one attempt in a given week. In Columns 4-9, the dependent variable is the number of attempts in the week. An *attempt* is defined as the first shipment between an exporter and a buyer. Columns 7-9 are estimated using a Poisson Pseudo-Maximum Likelihood (PPML) estimator. Control variables vary across specifications and include the log quantity and log value of the exporters' transactions in the week. Standard errors, in parentheses, are clustered at the exporter level. Asterisks next to the estimate, *, **, ***, denote statistical significance at the 0.10, 0.05, and 0.01 level, respectively.

D Additional Figures

Figure D.1: Average Shipment Unit Price: Relationships and Auctions

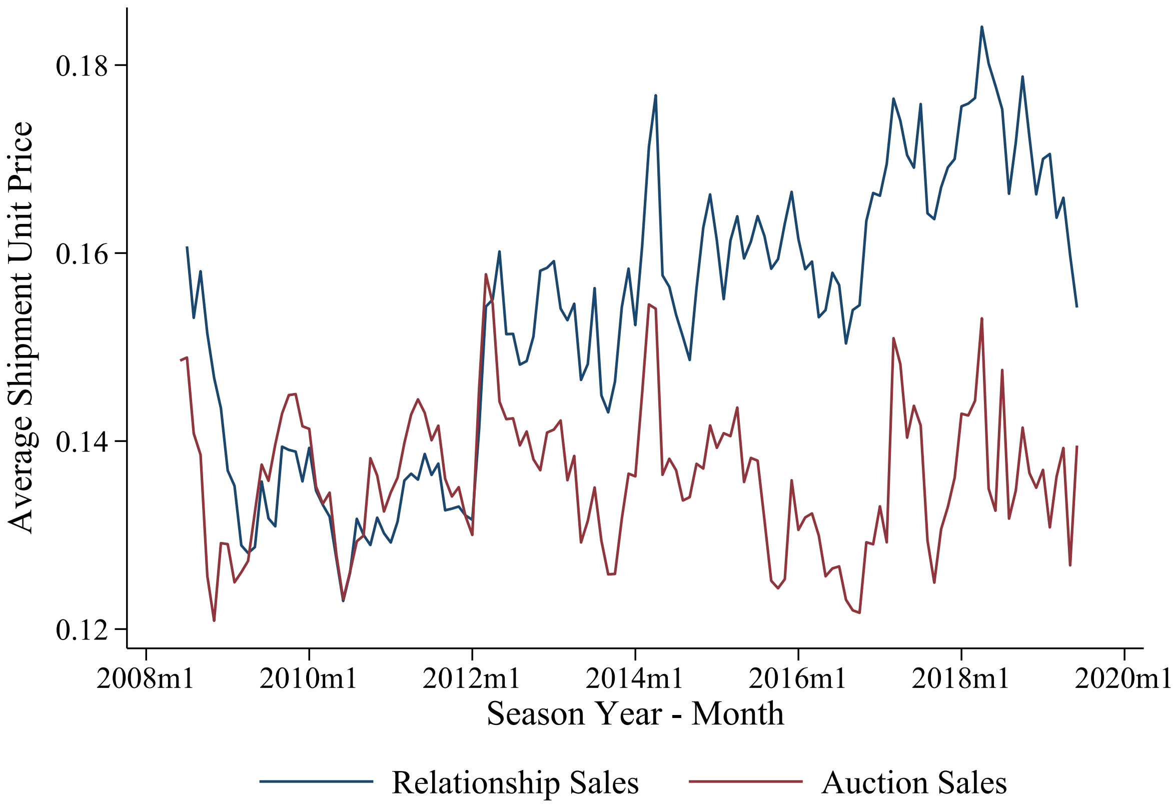

Note: The figure illustrates the monthly average transaction unit price for relationship sales and auctions, computed from customs data from 2009 to 2018. All prices are expressed in USD and averaged at the transaction level. Blue line denotes relationship sales, and red line denotes auction sales.

Figure D.2: Average Seasonal Exports by Ownership Type

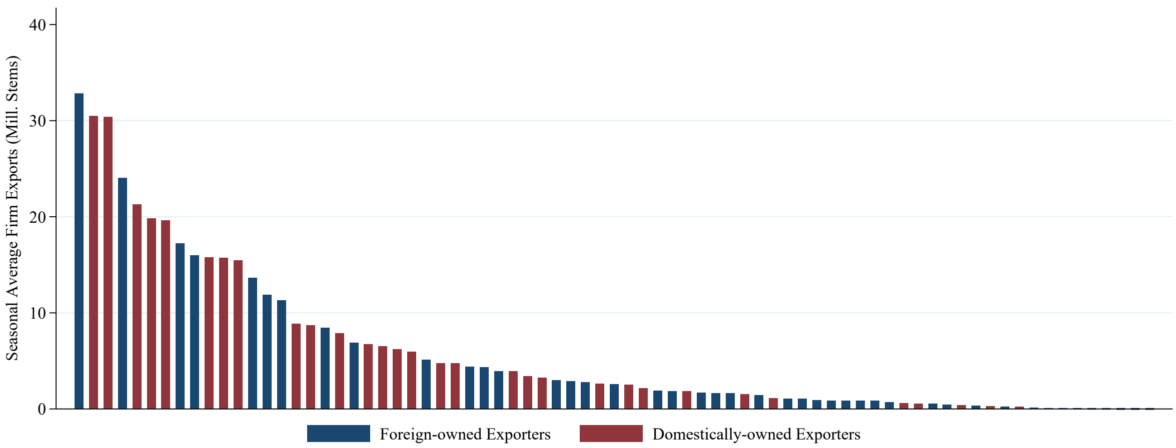

Note: The figure presents the seasonal average total firm exports during the flower season from 2009 to 2019, ranked from largest to smallest. Each firm is classified as either domestic or foreign-owned. Export values, expressed in millions of flower stems, are calculated from customs data. Blue bars represent foreign-owned exporters, while maroon bars represent domestically owned exporters.

References

Jovanovic, Boyan. 1979. "Job matching and the theory of turnover." *Journal of Political Economy*, 87(5): 972–990.

Ljungqvist, Lars, and Thomas J Sargent. 2018. Recursive Macroeconomic Theory. MIT press.

Nourdin, Ivan, and Giovanni Peccati. 2009. "Stein's method on Wiener chaos." *Probability Theory and Related Fields*, 145: 75–118.