Paying With Promises

Mitchell Watt*

July 2026

Abstract

Wages are often fixed even though tasks vary in difficulty. I study how a firm motivates workers when task assignment is its only incentive instrument. Workers observe task costs and may refuse or quit; the firm commits to compensating unpleasant work with better future assignments. With one worker, the optimal contract assigns tasks above a cutoff that rises with promised utility. With several workers, larger promises earn better tasks on average; workers complete every positive-value task and the least costly negative-value tasks. Pooling tasks among workers adds value, and spare workers expand capacity. Dynamics add value exactly when a static rule assigning all tasks is infeasible. Relative patience governs long-run assignments: negative tasks eventually disappear when workers are more patient, assignments become effectively stationary under equal patience, and difficult tasks recur when the firm is more patient. The model rationalizes task rotation and seniority rules as substitutes for task-contingent pay.

Keywords: Non-monetary incentives, dynamic incentives, task assignment, dynamic contracts, promised utility.

JEL Codes: C73, D82, D86

---

*Department of Economics, Monash Business School, 900 Dandenong Rd, Caulfield East, Victoria, Australia. Email: mitch.watt@monash.edu. Thank you to Paul Milgrom, Ravi Jagadeesan, and seminar participants at Stanford University and Monash University for helpful advice and comments related to this project.

1 Introduction

Many jobs comprise tasks with widely varying difficulty—dealing with different kinds of customers, working in irregular locations, or performing demanding versus routine work. Yet wages are often fixed at an annual or hourly rate and do not adjust task by task. A wage that compensates workers for the average task necessarily underpays them for difficult ones. Workers may respond by refusing an assignment, and they can always quit. How does a firm incentivize workers to exert costly effort when it cannot pay a premium for it?

In practice, many organizations compensate difficult assignments with better ones later. Rideshare platforms, whose fares do not price every difference between trips, let drivers who accept short trips from the airport keep their place at the front of the airport queue. Queensland's education department awards teachers transfer points for service in remote communities, which convert into priority over future postings. The U.S. Navy offsets demanding sea tours with subsequent shore duty. These schemes track who has borne the brunt of difficult work and prioritize them for later assignments. This paper develops a theory of such arrangements. I characterize the optimal dynamic contract for a firm whose only incentive instrument is task assignment, and show that the optimum has exactly the structure these institutions approximate: by accepting difficult tasks, workers earn a claim to future assignments, which the firm settles by delivering desirable ones.

Why don't firms simply price each task? Managers report that adjusting pay to circumstances violates fairness norms and damages morale (Akerlof & Yellen, 1990; Bewley, 1999). When performance is multidimensional or hard to contract on, task-contingent pay may invite gaming and organizations can instead rely on nonpay instruments (Baker, Gibbons, & Murphy, 1994; Holmström, 1984). I therefore take wage rigidity as given and study the instrument that remains: whether, when, and to whom tasks are assigned.

I formulate a dynamic principal–agent model in which a firm allocates tasks to an employee. Each period a task arrives whose net value to the worker—the wage less a random effort cost—may be positive or negative and is observed by both parties. The firm benefits from every completed task, so it wants all tasks done; the worker wants to do only the easier ones. The worker can refuse an assignment and quit at any time, so the contract must satisfy an *ex post participation constraint*: after seeing today's task and what the contract promises for the future, the worker must prefer completing the task to quitting. The firm commits to both on- and off-path continuation rules: it honors promised favorable assignments and permanently excludes a worker after refusal. Task values are public, so the problem is one of participation rather than screening. With task-contingent transfers, the firm would assign every task with positive total surplus and compensate the worker in each period. Without such transfers, compensation is paid in kind and in arrears by determining the tasks faced by the worker tomorrow. The firm and the worker may discount the future at different rates, and the ratio of their discount factors pins down the long-run dynamics of the employment relationship.

The optimal contract is characterized by a single state variable: the worker's promised utility, which records the future assignment value owed to the worker. When the firm assigns a task the worker would rather refuse, it must provide enough continuation utility to make acceptance worthwhile. Desirable tasks deliver current utility toward what the firm owes and reduce the continuation utility otherwise required. Characterizing the optimal contract involves calculating the shadow cost of promises. Small promises cost the firm nothing: it can honor them with tasks it wants to assign anyway (because workers and firms both agree they want to complete positive-value tasks). Larger promises require shielding the worker from tasks the firm would like done, so the firm's value falls and each additional unit of promised utility costs more than the last.

Section 3 gives a characterization of the optimal contract with a single worker. I first identify when a dynamic contract can *strictly* improve over a static contract. A static contract is optimal if and only if the average task is attractive enough for the worker to accept even the worst realized task, in which case the firm assigns all tasks. Otherwise, every optimal contract conditions on history (Proposition 4). A contract (Theorem 1) consists of an assignment rule and an updating rule for the promised utility. The optimal assignment rule is a cutoff rule: a task is assigned when its value exceeds a threshold depending on the current promised utility. The updating rule is built around a *target* for next period's promised utility: a promise the firm chooses whenever the worker's participation constraint is slack, chosen to smooth the marginal cost of promises over time. That target may lie above or below the current promised utility. It is above when the worker is more patient than the firm, below when the firm is more patient. Tasks fall into three regions. Favorable tasks are assigned, and the promised utility moves to the target. Moderately unpleasant tasks are assigned, but acceptance requires more than the target: the firm raises the promised utility exactly enough to secure acceptance. Sufficiently unpleasant tasks are rejected—the flow benefit of completing the task is insufficient to cover the cost of the minimal promise needed to incentivize acceptance—and the promised utility again moves to the target. The cutoff rule is weakly increasing in promised utility: the more the firm already owes a worker, the more selective it becomes about burdening the worker further.

This characterization delivers sharp comparative statics. The firm's payoff rises with the worker's patience, because a patient worker values future favors more and acceptance becomes cheaper to buy: as worker patience vanishes, the firm can assign only pleasant tasks, while as it approaches one, the firm can assign all tasks. First-order stochastic improvements in the task distribution likewise raise the firm's payoff. Firm patience cuts the other way, and in a perhaps surprising direction: a more patient firm earns a strictly *lower* normalized payoff, because it places more weight on later periods in which it is repaying promises created by earlier assignments. Yet those same promises sustain the completion of difficult tasks. Once the firm is at least as patient as its workers, the fraction of tasks completed stays strictly above the positive-task rate in the long run; when the firm is less patient, that fraction converges to the positive-task rate. The gain from dynamic contracting over the static benchmark is accordingly largest for impatient firms.

The long-run performance of the employment contract depends on who is more patient (Proposition 5). If the worker is more patient than the firm, the firm prefers to defer compensation: postponing a reward costs the patient worker little and benefits the impatient firm, so the optimal contract backloads incentives. The worker's promised utility ratchets upward, and negative assignments cease after a finite date—the worker is eventually tenured. If instead the firm is more patient, it prefers to pay down what it owes quickly: it reduces the promised utility whenever the worker's participation is not at issue, while difficult assignments recur with positive long-run frequency, pushing the promised utility back up. The result is a recurrent pattern in which the worker builds up his promised utility by accepting hard tasks and is repaid through stretches of favorable treatment. Under equal patience there is no systematic pull toward either boundary: promised utility converges, and the contract approaches a stationary cutoff.

Section 4 extends the model to N workers and $K \leq N$ tasks per period, allocated by a matching. The single-worker analysis asked only *when* tasks should be assigned; with several workers, the firm must also decide *who* receives each assigned task and how outstanding promised utility claims should be distributed across the workforce. A first question is whether workforce size and matching flexibility matter at all (Proposition 7). With a single task per period they do not: promised utilities aggregate, so any number of workers is equivalent to one worker with the same total promised utility. On the other hand, pooling multiple tasks creates value: assigning $K \geq 2$ tasks jointly strictly outperforms K separate single-worker relationships, because the firm can steer desirable tasks toward workers whose claims are costly to honor and unpleasant ones toward workers with small promises. More than K employees can also be strictly beneficial: an additional worker provides spare capacity, preserving the firm's ability to assign tasks while a colleague who accepted an exceptionally unpleasant assignment is being repaid with favorable treatment.

In the multi-worker model, the principal's value function is Schur-concave: for a fixed total promise, the firm prefers to divide it more evenly among workers. As a result of this equalization motive, an optimal matching (Theorem 2) assigns workers with larger promised utility higher-valued tasks *on average*. When the marginal costs of workers' promises differ, higher-value tasks are always assigned to workers with higher promised utility. When the marginal costs are tied—which can occur when their promises are close—the higher-promise worker sometimes receives a lower-value task. Promised utility thus functions as a priority score: accepting difficult work can raise a worker's rank, while desirable tasks flow, on average, to the highest-ranked. Every positive-value task is assigned, and negative-value tasks are assigned from the least bad downward, up to an endogenous stopping point.

The long-run dynamics of the optimal contract carry over from the single-agent model (Theorem 3). When workers are more patient than the firm, negative-value tasks are assigned only finitely often: beyond some date, the firm assigns exactly the pleasant tasks and nothing else. Under equal patience, the assignment rule becomes effectively stationary, while positive and negative assignments continue with positive long-run frequency. When the firm is more patient, there is an optimal contract whose assignment priorities follow a recurring pattern: at a positive fraction of dates, every task is undesirable and assigned, and the worker who takes the worst task moves from the lowest to the highest priority in the next period.

The paper makes three contributions relative to the literature reviewed below. First, it characterizes the optimal dynamic contract when task assignment is the only instrument and the firm faces an *ex post* participation constraint in each period. Second, by allowing the firm and its workers to differ in patience, it identifies relative patience as a key determinant of the long-run character of the workplace: careers that culminate in protected seniority, stationary assignment rules, or recurrent patterns of effort and reward. This margin is invisible in related models with a common discount factor (Forand & Zápal, 2020). Third, the multi-worker analysis shows that priority systems arise endogenously: whereas pseudo-market mechanisms endow agents with artificial budgets by design (Casella, 2005; Hylland & Zeckhauser, 1979), here the currency—priority for future tasks—emerges as the optimal contract's own state variable. This provides a foundation for the transfer points, priority queues, and rotation schedules that organizations actually use, which I revisit in Section 5.

The remainder of this section reviews the related literature. Section 2 introduces the model. Section 3 analyzes the single-worker problem and Section 4 the multi-worker problem. Section 5 maps the theory to organizational practice, and Section 6 concludes. Proofs are collected in the appendix.

1.1 Related Literature

The closest paper is Forand and Zápal (2020), who study repeated project opportunities with nontransferable utility and an agent who can leave the relationship. They show that production follows a cost-effectiveness ordering and that the contract becomes progressively more favorable to the agent. Their agent chooses whether to remain after observing the project but before a public randomization determines production. Here the worker observes the realized assignment and its continuation promise before deciding whether to accept. Participation must therefore hold for each realized assignment, which creates a task-specific floor on continuation utility. The task structure then delivers an explicit assignment cutoff and a target for continuation utility. Allowing the firm and the worker to discount at different rates shows when compensation is backloaded, stationary, or recurrent. The extension to several workers and simultaneous tasks also yields endogenous assignment priorities, gains from pooling task streams, and a distinct role for spare workforce capacity. Bird and Frug (2022) place the gradual generosity found by Forand and Zápal (2020) within a general theory of activities that become more favorable to the agent in long-term contracts with a common discount factor. This paper characterizes the allocation and continuation policy directly. Unequal discounting then separates eventual protection from unpleasant tasks, stationary assignment, and recurrent patterns of difficult work and compensation.

Several papers study dynamic job or task assignment inside organizations. Rotation can manage private information about workers or projects (de Clippel, Eliaz, Fershtman, & Rozen, 2021; Prescott & Townsend, 2006), while task assignment can be combined with training, wages, or relational incentives (Andrews & Barron, 2016; Bird & Frug, 2021; Fudenberg & Rayo, 2019). The most closely related multi-worker models use assignment itself to provide incentives. Ji and Vong (2025) study repeated assignment of a single task when effort is hidden and characterize first-best equilibria with a strict priority ranking that evolves with observed output. Eliaz, Fershtman, and Frug (2025) study optimal scheduling when workers privately observe task completion and idle time is the available reward. Their optimal schedule alternates between efficient production and delay. The present paper takes task values and actions as public and makes acceptance of a realized assignment the operative constraint. The firm commits to a payoff-maximizing contract and may face several heterogeneous tasks at once. These features produce value-based task selection, matching across workers, and priorities that can be soft when workers have equal marginal promise costs. They also separate the value of pooling tasks from the value of expanding the workforce. In the rideshare application, Castro, Ma, Nazerzadeh, and Yan (2021) design queue rules that use waiting time to reduce drivers' rejection of undesirable trips. This paper derives assignment priority from the optimal employment contract and characterizes how accepted assignments change it over time.

The paper also contributes to the literature on dynamic incentive provision without monetary transfers. Bird and Frug (2019) study privately observed investment and reward opportunities and show that incentives can be provided by giving the agent a predetermined interval during which every reward opportunity may be taken. Frankel (2016) derive contracts that constrain the discounted frequency of actions, while Guo and Hörner (2018) and Balseiro, Gurkan, and Sun (2019) use future allocation rights to elicit private values. These papers address information revelation. Task values are public here, and promised utility records the future assignment value needed to secure acceptance of assigned work. Promised utility is an endogenous liability created by the contract; artificial accounts in reporting mechanisms serve a different incentive purpose. The continuation-value logic also resembles favor exchange, in which costly actions today are compensated by favorable treatment later (Abdulkadiroğlu & Bagwell, 2013; Samuelson & Stacchetti, 2017). Those relationships are sustained bilaterally through future cooperation. Here a committed firm controls the task stream; workers may refuse and leave; the firm repays obligations generated by difficult work.

A related allocation literature studies how to distribute goods, positions, or unpleasant tasks without payments. Static mechanisms use artificial budgets, linked decisions, or random assignment to relax incentive and efficiency constraints (Casella, 2005; Goldlücke & Tröger, 2018; Hylland & Zeckhauser, 1979; Jackson & Sonnenschein, 2007). In a large dynamic assignment market, Combe, Nora, and Tercieux (2025) obtain spot mechanisms with virtual prices and artificial budgets carried across dates. This paper studies a finite employment relationship in which an assignment determines both production and compensation and must satisfy ex post participation. The relevant balance is derived from the contracting problem: difficult work raises promised utility, and favorable assignments redeem it. This provides a foundation for the point balances, queue positions, and priority rights used in practice.

Finally, Krasikov, Lamba, and Mettral (2023) study unequal discounting in a dynamic screening problem with persistent private information, transfers, and limited commitment. They show that the intertemporal cost of incentives can generate cycles with infinite memory. The public-information environment here isolates the cost of shifting compensation in kind across time. Starting from no outstanding promises, relative patience produces a three-way classification: unpleasant tasks eventually disappear when workers are more patient, assignment becomes effectively stationary under equal patience, and difficult work recurs when the firm is more patient. Methodologically, the analysis uses promised utilities and recursive methods from dynamic contracting and repeated games (Abreu, Pearce, & Stacchetti, 1990; Spear & Srivastava, 1987). It applies those methods to derive explicit rules for task selection, matching across workers, and the evolution of workplace priorities.

2 Model

In this section, I introduce a simple model of task assignment in a firm. Workers are ex ante symmetric, there is no private information, and all actions are publicly observed. The *only* contracting friction is that wages are fixed, while each worker can refuse an assignment and quit the job in any period. I formulate the firm's problem as an optimal contracting problem with an ex post participation constraint: after seeing the current task and promised continuation, each worker must prefer accepting the assignment to leaving.

Environment and timing. A principal employs $N \geq 1$ risk-neutral workers and faces $K \leq N$ indivisible tasks in each period. Workers are indexed by $i \in \{1, \dots, N\}$, tasks by $k \in \{1, \dots, K\}$, and dates by $t = 0, 1, \dots$. Each worker can complete at most one task per period, and each task can be assigned to at most one worker.

Task k at date t has effort cost $\kappa_{k,t}$. Every worker receives the same exogenous wage w, so the task's net value to any worker is $V_{k,t} := w - \kappa_{k,t}$. The values are independently and identically distributed across tasks and dates with distribution F, which is supported on a compact interval $\mathcal{V} = [\underline{v}, \overline{v}]$, where $\underline{v} < 0 < \overline{v}$, and has a continuous, positive, bounded density f. I write $V \sim F$ and $v \in \mathcal{V}$ for a generic task value and its realization. I write $\mathbf{V} = (V_1, \ldots, V_K) \sim F^K$ and $\mathbf{v} = (v_1, \ldots, v_K)$ for the corresponding task vector and its realization.

At the beginning of a period, the contract and history of assignments to the workers are publicly known. The task vector is then drawn and publicly observed. The principal proposes a matching, and the contract specifies the continuation following every vector of acceptance and refusal decisions. Each assigned worker observes her task and these continuations before deciding whether to accept. Accepted tasks are completed; rejected tasks are not. All task values, assignments, and acceptance decisions are observed by everyone.

Let $x_{ik,t} = 1$ if worker i completes task k at date t, and set it to zero otherwise. Feasibility requires $\sum_i x_{ik,t} \le 1$ for every task and $\sum_k x_{ik,t} \le 1$ for every worker. The principal obtains one unit of flow payoff from each completed task. Worker i receives flow utility $V_{k,t}$ if she completes task k. With discount factors $\delta_P, \delta_A \in (0,1)$, their normalized payoffs are

$$U_P = (1-\delta_P) \sum_{t=0}^{\infty} \delta_P^t \sum_{i=1}^{N} \sum_{k=1}^{K} x_{ik,t}, \qquad U_{A,i} = (1-\delta_A) \sum_{t=0}^{\infty} \delta_A^t \sum_{k=1}^{K} V_{k,t} x_{ik,t}.$$

Workers may refuse any proposed assignment and may also quit at any date and obtain continuation utility zero. When indifferent, a worker accepts a task. The principal can commit to history-dependent assignment and continuation rules, including withholding desirable future tasks even when doing so reduces her short-run payoff. In particular, the principal can punish off-path actions by withholding all future assignments.

Contracts. Let $\mathcal{H}_t$ be the public history through the end of date $t-1$. A contract specifies, after every $(\mathbf{v}, h_t)$, a (possibly randomized) feasible matching and a continuation contract following every vector of acceptance and refusal decisions. The principal commits to the contract at date zero and chooses it to maximize her expected payoff, anticipating the workers' acceptance and quitting decisions.

Following Abreu, Pearce, and Stacchetti (1990), it suffices to represent each continuation contract by a vector u_t of promised utilities and to restrict attention to contracts that are Markov in u_t: conditional on the current task vector, the matching and continuation-promise rules depend on the public history only through u_t.

Promised utilities and their feasible set. A promise vector $u = (u_1, \ldots, u_N)$ guarantees each worker i expected normalized utility of at least u_i from the current date onward. I now determine the set of feasible promise vectors in any dynamic contract.

Write $V_{(1)} \leq \cdots \leq V_{(K)}$ for the order statistics of $\mathbf{V}$. For $m \in \{0, 1, \ldots, K\}$, define the largest expected utility that m workers can jointly receive in a period by

$$\overline{V}_m := \mathbb{E}\left[\sum_{r=K-m+1}^{K} [V_{(r)}]_+\right], \qquad \overline{V}_0 := 0.$$

For $m > K$, set $\overline{V}_m := \overline{V}_K$.

For any set of workers S, at most $\min\{|S|, K\}$ tasks can be assigned to its members. Their total current utility cannot exceed the sum of the corresponding number of best nonnegative task values. Discounting and taking expectations therefore gives the necessary bounds

$$\sum_{i \in S} u_i \leq \overline{V}_{|S|} \qquad \text{for every } S \subseteq \{1, \ldots, N\}.$$

In Appendix A.1, I prove that the inequalities above are jointly sufficient and characterize the feasible promise set:

$$\mathcal{U}^{N,K} := \left\{ u \in \mathbb{R}_+^N : \sum_{i \in S} u_i \leq \overline{V}_{|S|} \text{ for every } S \subseteq \{1, \ldots, N\} \right\}. \tag{1}$$

Equivalently, define

$$\bar{v} := \big(\mathbb{E}[[V_{(K)}]_+], \mathbb{E}[[V_{(K-1)}]_+], \ldots, \mathbb{E}[[V_{(1)}]_+], 0, \ldots, 0\big) \in \mathbb{R}^N.$$

Then $\mathcal{U}^{N,K} = \{u \in \mathbb{R}_+^N : u \preceq_w \bar{v}\}$, where $\preceq_w$ denotes weak majorization:

$$u \in \mathcal{U}^{N,K} \quad \Longleftrightarrow \quad \sum_{j=1}^{m} u_j^{\downarrow} \leq \sum_{j=1}^{m} \bar{v}_j = \overline{V}_m \qquad \text{for } m = 1, \ldots, N.$$

Here $u^{\downarrow}$ is u sorted in weakly decreasing order. The promise set is therefore a compact, convex, symmetric, and coordinatewise downward-closed polytope. In Appendix A.1, I prove that the promise set is also *self-generating*: every promise in $\mathcal{U}^{N,K}$ can be delivered by an assignment rule and continuation promises that also belong to $\mathcal{U}^{N,K}$.

Recursive formulation. Following Myerson (1982), workers' acceptance decisions can be incorporated into the assignment rule without loss. In particular, every feasible contract has a payoff-equivalent representation in which prescribed assignments are accepted on path and a worker who refuses receives no future assignments.[1] In this representation, $X_{ik} = 1$ means that worker i is assigned task k and completes it. The requirement that workers accept their prescribed assignments is the *participation constraint* (PC_{ik}). If a task has value v, this constraint requires $(1-\delta_A)v + \delta_A u_i' \geq 0$. Write $\ell(v) := [-(1-\delta_A)v/\delta_A]_+$ for the smallest continuation promise satisfying the constraint, which I call the *participation floor*.

The *promise-keeping constraint* requires the contract to deliver each worker at least the utility that the firm has promised. Thus, at state u, worker i's expected normalized current utility plus discounted continuation utility must be at least u_i.

A deterministic Markov decision consists of a matching rule $X : \mathcal{V}^K \times \mathcal{U}^{N,K} \to \{0,1\}^{N\times K}$ and a continuation-promise rule $u' : \mathcal{V}^K \times \mathcal{U}^{N,K} \to \mathcal{U}^{N,K}$. The value function $\Phi^{N,K} : \mathcal{U}^{N,K} \to \mathbb{R}$ satisfies

$$\Phi^{N,K}(u) = \max_{X,u'} \mathbb{E}\left[(1-\delta_P)\sum_{i=1}^{N}\sum_{k=1}^{K} X_{ik}(\mathbf{V};u) + \delta_P \Phi^{N,K}(u'(\mathbf{V};u))\right], \tag{MOPT}$$

$$\mathbb{E}\left[(1-\delta_A)\sum_{k=1}^{K} V_k X_{ik}(\mathbf{V};u) + \delta_A u_i'(\mathbf{V};u)\right] \geq u_i \qquad \text{for every } i, \tag{PK$_i$}$$

$$u_i'(\mathbf{v};u) \geq \ell(v_k) X_{ik}(\mathbf{v};u) \qquad \text{for every } i,k,\mathbf{v}, \tag{PC$_{ik}$}$$

$$\sum_{i=1}^{N} X_{ik}(\mathbf{v};u) \leq 1 \quad \text{and} \quad \sum_{k=1}^{K} X_{ik}(\mathbf{v};u) \leq 1 \qquad \text{for every } i,k,\mathbf{v}. \tag{MATCH}$$

The maximization is over measurable rules. Constraint (PK_i) is promise keeping, constraint (PC_{ik}) is the participation constraint, and constraint (MATCH) makes each realized allocation a matching. An *optimal Markov contract* is a measurable assignment and continuation rule that attains this maximum at every promise vector. In Appendix A.1, I prove that this maximum is attained and that allowing lotteries over one-period decisions does not increase the value. Thus an optimal contract can be chosen Markov and deterministic, with a single matching and continuation promise selected for each state and realized task vector.

Initial promises. Before date zero, worker i can decline the relationship and obtain reservation utility R_i. Assume $R = (R_1, \ldots, R_N) \in \mathcal{U}^{N,K}$. The principal starts the recursive contract at $u_0 = R$; the monotonicity results below show that she never benefits from promising more initially. The long-run results use the normalization $R = 0$, while the value-function and policy characterizations apply from every feasible promise.

[1]Starting from any contract satisfying the participation constraints, preserve the prescribed continuation for workers who accept and never assign another task to a worker who refuses. Removing a refusing worker relaxes the matching constraints, gives that worker continuation utility zero, and weakly strengthens acceptance incentives. Because all prescribed assignments are accepted on path, this changes only off-path play and leaves every equilibrium payoff unchanged.

3 The single-worker benchmark

The single-worker case isolates the central dynamic tradeoff in the model. With no matching decision, the firm chooses only whether to assign the realized task and how much future utility to offer in return for accepting an unpleasant one. This case shows when history dependence creates value and how the relative patience of the firm and worker shapes the evolution of the employment contract over time. It also provides a benchmark for separating these dynamic gains from the additional gains created by matching several tasks to several workers.

In this section, set $N = K = 1$ and write $V^+ := \mathbb{E}[V_+]$, so the feasible promise set is $\mathcal{U}^{1,1} = [0, V^+]$. I write $\Phi := \Phi^{1,1}$ for the value function.

Value function properties. The following result records the key properties of the value function that shape the optimal policy.

Proposition 1 (Single-agent value function properties)**.** *For $N = K = 1$, (MOPT) has a unique bounded fixed point $\Phi : [0, V^+] \to \mathbb{R}$, which is continuous, nonincreasing, and concave. Moreover, there exists $\bar{u} \in (0, V^+)$ such that Φ is constant on $[0, \bar{u}]$ and strictly concave on $(\bar{u}, V^+]$.*

Proposition 1 is proved in Appendix A.2. I call the interval $[0, \bar{u}]$ the *plateau*. The plateau reflects rents that the worker receives under the firm's preferred assignment policy. The firm always assigns positive-value tasks, which give the worker strictly positive utility, while (PC_{ik}) ensures that every assigned negative-value task is accompanied by enough continuation utility to offset its current cost. Consequently, even a contract starting from a zero promise delivers positive expected utility. Any promise small enough to be covered by this utility can be satisfied without changing assignments or reducing the firm's payoff. The promise-keeping constraint therefore remains slack on $[0, \bar{u}]$. Above $\bar{u}$, delivering additional utility requires the firm to shield the worker from some unpleasant tasks, and its value begins to fall. Figure 1 illustrates this geometry.

Static benchmark. Before characterizing the optimal dynamic contract, I first study what the firm can achieve with a static assignment rule. A *static contract* is an assignment rule $X : \mathcal{V} \to \{0, 1\}$ applied identically after every history. The associated stationary promise is $U_X := \mathbb{E}[vX(v)]$, the expected value the agent receives each period. Feasibility requires that each assigned task be worth accepting given the stationary continuation: $(1 - \delta_A)v + \delta_A U_X \geq 0$ for every v with $X(v) = 1$. The principal's payoff is $\Pi(X) := \mathbb{E}[X(v)]$.

Since each completed task yields the same unit payoff to the firm regardless of v, while the compensation required for acceptance is decreasing in v, the firm strictly prefers to assign the highest-value tasks. Cutoff rules are therefore optimal without loss of generality.

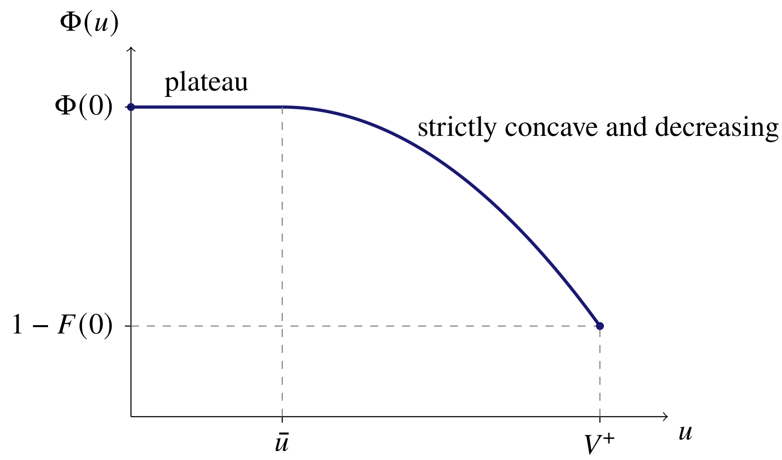

Figure 1: The shape of the single-worker value function.

For a candidate cutoff $c < 0$, define $U(c) := \mathbb{E}[v\mathbf{1}\{v \geq c\}]$, the stationary promise generated by the cutoff rule. The marginal assigned task requires compensation $\frac{1-\delta_A}{\delta_A}(-c)$. Thus $H(c) := U(c) + \frac{1-\delta_A}{\delta_A}c \geq 0$ is exactly the condition under which the cutoff c is statically feasible. The principal seeks to push the cutoff as low as possible subject to feasibility, so the threshold $H(c) = 0$ identifies the limit. The following is proved formally in Appendix A.3.

Proposition 2 (Optimal static contract)**.** *Let*

$$c^S := \begin{cases} \underline{v}, & \text{if } H(\underline{v}) \geq 0, \\ \text{the unique root of } H(c) = 0 \text{ in } (\underline{v}, 0), & \text{if } H(\underline{v}) < 0. \end{cases}$$

Then the cutoff rule $X^S(v) = \mathbf{1}\{v \geq c^S\}$ *is an optimal static deterministic contract. Its stationary promise is* $U^S = \mathbb{E}[v\mathbf{1}\{v \geq c^S\}]$*, and the principal's payoff is* $\Pi^S = \Pr(v \geq c^S)$.

The static benchmark also gives transparent comparative statics in δ_A, δ_P and F.

Proposition 3 (Comparative statics of the static contract)**.** *As* δ_A *increases (holding* δ_P *and* F *fixed), the optimal static cutoff* c^S *weakly decreases, the assignment probability* Π^S *weakly increases, and the worker's stationary payoff* U^S *weakly decreases. These objects do not depend on* δ_P*. Holding the discount factors fixed, a first-order stochastic improvement in the task-value distribution weakly increases* Π^S.

A more patient worker places greater weight on the continuation reward and thus needs fewer favorable future tasks to accept an unpleasant task today. The firm uses this flexibility to lower the cutoff and assign more tasks, reducing the worker's stationary payoff. The firm's patience has no effect because a static rule generates the same assignment rate and worker payoff in every period. Under a first-order improvement in task values, the old assignment rule remains feasible and the firm can weakly increase the fraction assigned. The static cutoff and the worker's stationary payoff are not generally ordered: the firm may use better task values to add negative tasks.

Given a promised utility u, let $c^S(u) := \inf\{c \in [\underline{v}, 0] : U(c) \geq u \text{ and } \ell(c) \leq u\}$, with $c^S(V^+) = 0$. This is the lowest stationary cutoff that can deliver promise u while satisfying (PC_{ik}) for every assigned task.

When is the static contract dynamically optimal? When $\mathbb{E}[v] \geq \frac{1-\delta_A}{\delta_A}(-\underline{v})$, it is feasible for the firm to assign all tasks with a static assignment rule: the value of the average task is enough to incentivize the agent to accept even the worst possible task. In that case, the principal achieves payoff one every period, which is the principal's maximum payoff each period, so the static assignment rule is also dynamically optimal. Conversely, when this condition fails, the optimal static cutoff excludes some tasks. The following shows that the principal can always do strictly better with a dynamic policy.

Proposition 4 (Static full assignment or dynamic improvement)**.** *Either:*

(i) the optimal static contract has cutoff $c^S = \underline{v}$, so it assigns every task, or

(ii) the optimal static contract has cutoff $c^S > \underline{v}$, and the optimal dynamic contract strictly improves on the optimal static contract: $\Phi(0) > \Pi^S$.

Moreover, when $c^S > \underline{v}$, every optimal Markov contract starting on the plateau $[0, \bar{u}]$ exits the plateau in finite time almost surely.

The key implication is that a static policy is optimal only when the static optimum allocates all tasks. When $c^S > \underline{v}$, every static rule leaves a set of low-value tasks unassigned, and a dynamic rule can improve on this by using future task selection to compensate the worker for accepting some of those tasks today. A contract that only used promises on the plateau could move the promise within $[0, \bar{u}]$, but this movement would not change assignment incentives: promised utility has zero marginal cost throughout this interval, so every plateau state gives the same optimal cutoff. Such a contract is therefore payoff-equivalent to a static cutoff contract, contradicting $\Phi(0) > \Pi^S$. The result is proved formally in Appendix A.3.

In light of the above, I maintain the dynamic case henceforth, by assuming that

$$\mathbb{E}[v] < \frac{1-\delta_A}{\delta_A}(-\underline{v}). \tag{NS}$$

Optimal dynamic contract and long-run dynamics I use Lagrangian methods to obtain the following characterization of the optimal contract.

Theorem 1 (Single-agent optimal contract)**.** *Suppose* (NS) *holds, and write* $\mathcal{P} := [0, \bar{u}]$ *for the plateau. The optimal contract satisfies:*

(i) ***Assignment.*** *There is a cutoff function* $\gamma : [0, V^+] \to [\underline{v}, 0]$*, with* $\gamma \equiv \gamma_0$ *on* $\mathcal{P}$*, nondecreasing on* $(\bar{u}, V^+]$*, and* $\gamma(V^+) = 0$*, such that* $X(v; u) = \mathbf{1}\{v \geq \gamma(u)\}$ *except at task values occurring with probability zero.*

(ii) ***Continuation promises.*** *If* $u \in \mathcal{P}$*, then*

$$u'(v;u) \in \begin{cases} \mathcal{P} \cap [\ell(v)X(v;u), V^+], & \ell(v)X(v;u) \leq \bar{u}, \\ \{\ell(v)\}, & \ell(v)X(v;u) > \bar{u}, \end{cases}$$

with the continuation promises across task realizations chosen to satisfy promise keeping. If $u \in (\bar{u}, V^+)$*, there is a target continuation promise* $t(u) \in [\bar{u}, V^+)$ *such that every optimal decision uses the unique continuation rule*

$$u'(v;u) = \begin{cases} t(u), & v < \gamma(u), \\ \max\{\ell(v),\ t(u)\}, & v \geq \gamma(u). \end{cases}$$

At $u = V^+$*, feasibility pins down* $u'(v; V^+) = V^+$ *almost surely.*

Figure 2 illustrates the allocation and continuation promises for $u \in (\bar{u}, V^+)$.

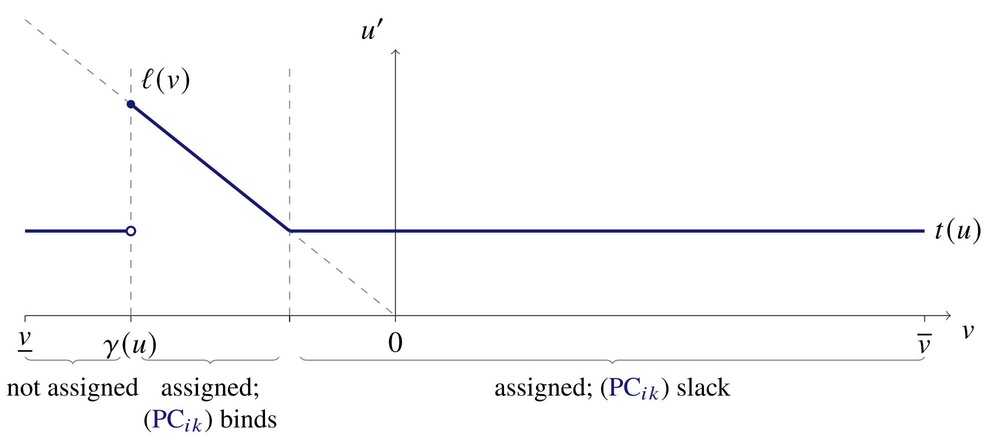

Figure 2: Assignment and continuation promises $u'(v;u)$ above the plateau. The navy line is the optimal continuation rule; the dashed diagonal is the continuation required for acceptance.

To explain the characterization, let $\lambda \geq 0$ be the Lagrange multiplier on the promise-keeping constraint. By the envelope theorem, $-\lambda$ is a supporting slope of Φ at u. Thus, define the *costate correspondence* by $\Lambda(u) := \partial(-\Phi)(u) \cap \mathbb{R}_+$: a costate $\lambda \in \Lambda(u)$ satisfies $\Phi(z) + \lambda z \leq \Phi(u) + \lambda u$ for every feasible promise z. Economically, λ is the marginal cost of raising promised utility at u. In the interior of the plateau $\mathcal{P} = [0, \bar{u}]$, the only costate is zero. The top state V^+ need not admit any finite costate.

Given λ, the principal compares the Lagrangian gain from assigning the realized task with the gain from leaving it unassigned. For a required continuation promise a, let $B_\lambda(a) := \max_{z \in [a, V^+]}\{\delta_P \Phi(z) + \delta_A \lambda z\}$, with $B_\lambda(a) = -\infty$ when $a > V^+$. After optimizing over continuation promises, the gain from assigning a task of value v is

$$\Delta(\lambda; v) := (1 - \delta_P) + (1 - \delta_A)\lambda v + B_\lambda(\ell(v)) - B_\lambda(0).$$

The first term is the firm's current payoff from completing the task, the second values the worker's current utility at its shadow cost, and the final difference is the cost of imposing the continuation promise required for acceptance. The task is assigned when $\Delta(\lambda; v) > 0$ and rejected when $\Delta(\lambda; v) < 0$; either decision is optimal at equality. This gain is increasing in v: higher values require less compensation and give the same unit payoff to the firm. Hence the optimal assignment decision is a cutoff in v, as in the static case.

The Lagrangian first-order conditions imply that at each date along an optimal path with $u_t < V^+$, a supporting costate $\lambda_t \in \Lambda(u_t)$ and a multiplier $\mu_t \geq 0$ on the participation constraint can be selected to satisfy the *costate equation*

$$\lambda_{t+1} = q\lambda_t + \frac{\mu_t}{\delta_P} \in \Lambda(u_{t+1}), \qquad q := \frac{\delta_A}{\delta_P}.$$

The zero costate supports every promise on the plateau, which is why there may be several possible continuation promise selections for $u \in \mathcal{P}$. Strict concavity above the plateau implies that each positive costate supports a unique promise. When (PC_{ik}) is slack, $\mu_t = 0$, so $q\lambda_t$ is the next-period costate. Above the plateau, the unique promise supported by this costate is the target continuation promise. If (PC_{ik}) requires a continuation above the target, the constraint binds and its multiplier raises the costate and the promise above this target. The target lies above the current promise when $q > 1$, equals it when $q = 1$, and lies below it when $q < 1$.

For this reason, q governs the contract's long-run dynamics. Under (NS), Proposition 4 implies that the promise process exits the plateau, resulting in a strictly positive costate, in finite time almost surely. Because $q > 0$ and $\mu_t \geq 0$, a positive costate remains positive; once the selected costate becomes positive, subsequent promises cannot enter $[0, \bar{u})$. Thereafter, q determines whether the costate process drifts up, stabilizes, or oscillates, and hence how promised utility evolves. To state a result on oscillation, I use the following definition. For a sequence $\{u_t\}$, the *omega-limit set* is

$$\omega(u_t) \;=\; \big\{x : u_{t_k} \to x \text{ for some subsequence } t_k \to \infty\big\},$$

the set of all subsequential limits of the path.

Proposition 5 (Single-agent long-run dynamics)**.** *Suppose* (NS) *holds. Let u_t be the promised continuation value of the agent in period t under any optimal Markov contract starting from $u_0 = 0$, and let γ denote its cutoff rule. Then:*

(i) ***Agent more patient (**$q > 1$**):*** *The process enters $(\bar{u}, V^+]$ in finite time almost surely. From that time onward, u_t is nondecreasing, strictly increases whenever $u_t < V^+$, and converges almost surely to V^+; hence $\gamma(u_t) \to \gamma(V^+) = 0$.*

(ii) ***Equal patience (**$q = 1$**):*** *The process enters $(\bar{u}, V^+)$ in finite time almost surely and never returns to $[0, \bar{u}]$. Promised utility converges almost surely to a history-dependent random limit $u_\infty \in (\bar{u}, V^+)$, and the cutoff converges almost surely to $\gamma(u_\infty)$. This limiting cutoff is the optimal stationary cutoff for delivering the limiting promise: $\gamma(u_\infty) = c^S(u_\infty)$ almost surely.*

(iii) ***Principal more patient (**$q < 1$**):*** *Let*

$$\gamma_0 := \gamma(0) = \gamma(\bar{u}) < 0, \qquad u^\sharp := \min\{V^+, \ell(\gamma_0)\} > \bar{u}.$$

The process satisfies $u_t \in [0, u^\sharp]$ for all t, enters $(\bar{u}, V^+]$ in finite time almost surely, and thereafter never falls below $\bar{u}$. Moreover, $\omega(u_t) = [\bar{u}, u^\sharp]$ almost surely.

When $q > 1$ (more patient agent), the next-period costate $q\lambda$ exceeds λ when the participation constraint is slack, so the target continuation promise lies above the current state and promised utility drifts upward monotonically after the process leaves the plateau. As the promise converges to V^+, the assignment cutoff converges to zero: negative-value tasks are eventually assigned with vanishing probability.

With equal patience ($q = 1$), promised utility is stationary when the participation constraint is slack and jumps upward when it binds. Thus, the promise process converges almost surely to a history-dependent limit $u_\infty \in (\bar{u}, V^+)$. The assignment rule converges to the static cutoff delivering this limiting promise.

When $q < 1$ (more patient principal), the next-period costate $q\lambda$ lies below λ when (PC_{ik}) is slack, so the target continuation promise is below the current state and promised utility drifts downward toward the plateau.

Negative assignments for which (PC_{ik}) binds periodically raise promised utility again, generating recurrent fluctuations on $[\bar{u}, u^\sharp]$.

Theorem 1 is proved in Appendix A.3, and Proposition 5 in Appendix A.5.4.

4 Optimal design with multiple workers and tasks

We now return to the firm with an arbitrary number of workers N and per-period tasks $K \leq N$. The central logic of the single-worker benchmark carries over, but the presence of several workers now adds a matching problem. After deciding which tasks to assign, the firm must decide who receives each one. It must also decide how to distribute outstanding promises across workers, because compensating one worker changes both her future assignments and the tasks available to everyone else. The firm can spread promises across workers, match tasks flexibly, and draw on workers who currently have smaller claims to favorable treatment. These possibilities create the new questions studied in this section: whether additional workers and flexible matching create value, how the firm distributes promises, which workers receive the better tasks, and how these assignments evolve over time. Symmetry and concavity lead the firm to prefer promises that are lower and more evenly distributed, providing the link between these questions.

Properties of the value function Characterizing $\Phi^{N,K}$ requires additional notation and definitions. As in Section 3, define the *costate correspondence* by $\Lambda^{N,K}(u) := \partial(-\Phi^{N,K})(u) \cap \mathbb{R}^N_+$; each coordinate λ_i measures the marginal cost of increasing worker i's promise. For $\lambda \in \mathbb{R}^N_+$, write $\Lambda^{-1}(\lambda) := \{u \in \mathcal{U}^{N,K} : \lambda \in \Lambda^{N,K}(u)\}$ for the promises at which λ is a valid vector of marginal costs. Write $\mathcal{P}^{N,K} := \Lambda^{-1}(0)$ for the value function's *plateau*, the set of promises at which the principal attains the same maximal value as at promise zero.

The single-agent value function has a useful property: once a costate is positive, that costate identifies a unique promise. With several agents, this is complicated by the possibility of ties in the costate vector; in that case, I show the *total* promise assigned to each group of workers with the same costate is unique, while the individual promises within such a group may remain undetermined. Formally, define $i \sim_\lambda j$ if and only if $\lambda_i = \lambda_j$. I call the resulting sets $B_1, \ldots, B_r$ the *equal-costate groups* of λ: each contains the workers whose promises have the same marginal cost. Write

$$S_\lambda z := \left(\sum_{i \in B_1} z_i, \ldots, \sum_{i \in B_r} z_i \right)$$

for the vector of total promised utilities for each equal-costate group.

Proposition 6 (Multi-agent value function properties)**.** *The Bellman equation* (MOPT) *has a unique bounded fixed point $\Phi^{N,K} : \mathcal{U}^{N,K} \to \mathbb{R}$, with the following properties:*

(i) *$\Phi^{N,K}$ is concave, continuous, coordinatewise nonincreasing, symmetric, and Schur-concave.*[2]

(ii) *$\mathcal{P}^{N,K} = \{u \in \mathcal{U}^{N,K} : \Phi^{N,K}(u) = \Phi^{N,K}(0)\}$ is nonempty, compact, convex, symmetric, and downward closed: if $u \in \mathcal{P}^{N,K}$ and $0 \le u' \le u$, then $u' \in \mathcal{P}^{N,K}$. It is a strict subset of $\mathcal{U}^{N,K}$ containing a nonzero promise vector.*

(iii) *For every strictly positive costate $\lambda \in \mathbb{R}^N_{++}$, the set $S_\lambda \Lambda^{-1}(\lambda)$ is singleton. In particular, if λ has no ties, then $\Lambda^{-1}(\lambda)$ contains exactly one promise vector.*

Proposition 6 is proved in Appendix A.2. It is the multi-agent analogue of the flat then curved shape of the single-agent value function. Concavity and symmetry imply Schur-concavity: for any fixed total level of promised utility, the principal prefers a more equal distribution of promises across agents. This equalization force is central to the characterization of the assignment below. The result on the inverse costate map is weaker than the strict concavity above the plateau obtained in the single-agent model. In the multi-agent model, a strictly positive costate vector pins down only the *total* promise in each equal-costate group; if its coordinates are distinct, it pins down the entire promise vector.

Workforce size and the benefits of pooling tasks. A first question is whether a firm that pools K workers and K task streams obtains more value than K separate single-worker contracts. A second question is whether the firm benefits from hiring more workers than per-period tasks.

Let $\phi := \Phi^{1,1}$ denote the value of the optimal single-worker contract given distribution F. The benchmark $K\phi(0)$ is therefore the payoff from using K copies of the optimal single-worker contract, with each worker permanently paired with one stream of tasks.

Proposition 7 (Workforce size and task pooling)**.** *The following comparisons hold.*

(i) *With one task per period, any number of workers is equivalent to one: for every $N \ge 1$ and $u \in \mathcal{U}^{N,1}$,*

$$\Phi^{N,1}(u) = \phi\left(\sum_{i=1}^{N} u_i\right).$$

(ii) *For every $K \ge 1$ and $N \ge K$,*

$$\Phi^{N,K}(0) \ge \Phi^{K,K}(0) \ge K\phi(0).$$

[2] A function W on a symmetric set is *Schur-concave* if $W(x) \le W(y)$ when x majorizes y.

(iii) Suppose $K \geq 2$. If static full assignment fails, so that $\mathbb{E}[V] < \ell(\underline{v})$, then pooling tasks adds value:

$$\Phi^{K,K}(0) > K\phi(0).$$

(iv) For every $K \geq 2$, there are discount factors and task-value distributions for which:

$$\Phi^{K+1,K}(0) > \Phi^{K,K}(0).$$

With one task per period, workers are perfect substitutes: only the total outstanding promise matters, so additional workers create no value. With several tasks, pooling task streams together allows the firm to direct better tasks toward workers whose promises are more costly to honor and less desirable tasks toward workers with smaller claims. This matching flexibility strictly improves on permanent worker–task relationships whenever those relationships cannot assign every task.

An additional worker can also provide spare capacity. A worker who accepts an exceptionally unpleasant task may need favorable treatment afterward, leaving fewer workers available for other tasks. A larger workforce can preserve assignment capacity while that promise is being honored. Part (iv) establishes that this gain arises in certain parametric examples. The proof is in Appendix A.6.

Static optimality and the plateau. Before analyzing the optimal dynamic contract, I first investigate when static contracts are sufficient. A static contract is a fixed promise vector $u^S \in \mathcal{U}^{N,K}$ together with a stationary (but possibly random[3]) assignment rule $X^S : \mathcal{V}^K \to [0,1]^{N\times K}$. The assignment $X^S(\mathbf{v})$ must satisfy $\sum_i X^S_{ik}(\mathbf{v}) \leq 1$ for every task k and $\sum_k X^S_{ik}(\mathbf{v}) \leq 1$ for every agent i. Constraint (PC$_{ik}$) requires $X^S_{ik}(\mathbf{v}) = 0$ whenever $u^S_i < \ell(v_k)$. Promise keeping requires $u^S_i \leq \mathbb{E}[\sum_k V_k X^S_{ik}(\mathbf{V})]$ for every agent i. The static payoff is $\Pi(X^S) := \mathbb{E}[\sum_{i,k} X^S_{ik}(\mathbf{V})]$, and Π^S denotes the largest payoff among feasible static contracts.

A static rule assigning all K tasks every period is clearly the optimal dynamic contract whenever the average value of a task is enough to incentivize an agent to accept the worst possible task. In that case, the firm attains the maximum payoff in every period.

Proposition 8 (Static optimality and dynamic improvement)**.** *If $\mathbb{E}[V] \geq \ell(\underline{v})$, a static full-assignment contract is feasible and attains the dynamic upper bound K for every $\delta_P \in (0,1)$. If $\mathbb{E}[V] < \ell(\underline{v})$, then the optimal dynamic contract* strictly *improves on the optimal static contract: $\Phi^{N,K}(0) > \Pi^S$. Consequently, a feasible*

[3] I allow randomization in the static benchmark because static contracts cannot use history-dependent promises to balance expected utility across agents. In the dynamic problem, continuation promises provide that balancing instrument, and deterministic Markov rules attain the same value as rules that randomize.

static contract is dynamically optimal from promise zero if and only if it assigns all K tasks after F^K-almost every realization of values.

Proposition 8 is proved in Appendix A.4. It first identifies the condition under which the static assignment $x_{ik} = 1\{i = k\}$, for $k = 1, \dots, K$, is feasible and shows that this is necessary and sufficient for all tasks to be assigned under a static contract. When the average task value is low, the optimal static contract leaves some tasks unassigned. A dynamic contract can do better by assigning mildly worse tasks at least sometimes when they appear, raising future promises for the workers who take them.

Characterizing the optimal assignment. I now characterize the optimal multi-worker task allocation. Say that a worker is *idle* in a period if she is assigned no task. To compare every worker's outcome under the optimal contract, represent idleness by a dummy task with value zero. Let $Y_i(\mathbf{v}; u)$ be the value of the actual or dummy task assigned to worker i under the matching prescribed at state u after task vector $\mathbf{v}$, and write $\bar{Y}_i(u) := \mathbb{E}[Y_i(\mathbf{V}; u)]$ for her expected current task utility.

Theorem 2 (Optimal task allocation)**.** *For every initial promise $u_0 \in \mathcal{U}^{N,K}$, there exists an optimal deterministic Markov contract with the following properties at every promise state u. If $u_i < u_j$, then $\bar{Y}_i(u) \leq \bar{Y}_j(u)$. After almost every task realization, the contract assigns every positive-value task and, if it assigns a negative-value task, it also assigns every higher-value negative task.*

Theorem 2 says that promised utility acts as a priority score. The optimal allocation is assortative in the sense that a worker with a larger outstanding promise receives a weakly higher task value on average. The second conclusion is the multi-worker analogue of the single-worker cutoff. The firm assigns every desirable task and then works down through the negative tasks from the least unpleasant, stopping when the additional task is no longer worth the additional continuation utility required for acceptance.

The proof begins by ranking workers according to the marginal cost of fulfilling their promises. Let λ_i denote the cost to the firm of marginally increasing worker i's promised utility. An exchange argument shows that, if $\lambda_i < \lambda_j$, worker j receives a weakly higher task value after almost every task realization. Intuitively, a desirable task delivers current utility and therefore helps the firm meet the promise that is more costly to provide through future assignments. Because the value function is Schur-concave, the firm weakly prefers promises to be distributed more equally across workers. By the Schur–Ostrowski criterion, a worker with a higher promise therefore cannot have a strictly lower marginal cost. Hence, when marginal costs differ, the worker with the higher promise receives a weakly higher task value after almost every task realization.

A complication arises when several workers have the same marginal cost. In that case, the firm values utility delivered to them at the same rate, so its one-period objective depends on their delivered utility only through the group total. Redistributing a fixed total promise within the group leaves the firm's value unchanged. Symmetry and convexity imply that every such region contains the point at which the members' promises are equal. It is therefore useful to think of promises that are sufficiently close to remain in the same tied-costate region as generating *soft priority*. Higher-promise workers may receive worse tasks in some realizations, but the optimal assignment described in the theorem gives them weakly greater *expected* current task utility. Once the promises are sufficiently different to generate distinct marginal costs, priority becomes strict and assignments must be assortative task by task.

The reason for soft priority when promised utilities are close is the promise-keeping constraint. When promises are close, *always* giving the better task to the higher-promise worker can create a difference in current utility that exceeds the difference in total promised utility. When promise keeping binds, $u_j - u_i = (1 - \delta_A)\mathbb{E}[Y_j - Y_i] + \delta_A \mathbb{E}[u'_j - u'_i]$. An excessive current advantage for worker j must then be offset by a continuation promise favoring worker i, which may be infeasible or costly. Occasional reversals in current assignments keep the average utility difference consistent with the promises while preserving optimality. The formal proof is in Appendix A.4.

A full characterization of a deterministic Markov contract also requires a rule for continuation promises. The cleanest description uses costates, the marginal costs of promised utility. Fix an optimal deterministic Markov contract. Along any realized path, the current and next costates $\lambda_t \in \Lambda^{N,K}(u_t)$ and $\lambda_{t+1} \in \Lambda^{N,K}(u_{t+1})$, together with multipliers μ_t on (PC$_{ik}$), can be chosen consistently so that, except on task realizations of probability zero at each history, the *costate equation* $\lambda_{t+1} = q\lambda_t + \mu_t/\delta_P$ holds, where $q = \delta_A/\delta_P$. A multiplier can be positive only when (PC$_{ik}$) binds. Thus relative patience scales the selected costates by q, while binding participation constraints can push them upward. In Appendix A.4, Proposition 10 and Lemma 12 derive this equation and construct a measurable sequence of consistent costates. When the selected next costate is strictly positive, Proposition 6 pins down the total continuation promise in each of its equal-costate groups.

Comparative statics. I now present some comparative statics of the optimal contract. These comparative statics apply to both the single- and multi-worker model. Write $P^{N,K}(\delta_A, \delta_P, F)$ for the firm's optimal expected payoff from an initial promise of zero as a function of δ_A, δ_P and F.

Proposition 9 (Multi-agent comparative statics)**.** *Fix $N \geq K$. The following comparisons hold.*

(i) ***Worker patience.*** *Fix δ_P and F. If $0 < \delta'_A < \delta''_A < 1$, then $P^{N,K}(\delta''_A, \delta_P, F) \geq P^{N,K}(\delta'_A, \delta_P, F)$, with strict inequality whenever $P^{N,K}(\delta'_A, \delta_P, F) < K$. Moreover,*

$$\lim_{\delta_A \downarrow 0} P^{N,K}(\delta_A, \delta_P, F) = K \Pr(V > 0), \qquad \lim_{\delta_A \uparrow 1} P^{N,K}(\delta_A, \delta_P, F) = K.$$

(ii) ***Firm patience.*** *Fix* δ_A *and* F*. If* $0 < \delta'_P < \delta''_P < 1$*, then* $P^{N,K}(\delta_A, \delta'_P, F) \geq P^{N,K}(\delta_A, \delta''_P, F)$*, with strict inequality when static full assignment is infeasible. By contrast, the unnormalized discounted task count is strictly increasing:*

$$\frac{P^{N,K}(\delta_A, \delta'_P, F)}{1 - \delta'_P} < \frac{P^{N,K}(\delta_A, \delta''_P, F)}{1 - \delta''_P}.$$

Because the optimal static payoff is independent of δ_P*, the gain* $P^{N,K}(\delta_A, \delta_P, F) - \Pi^S$ *from using a dynamic rather than a static contract is strictly decreasing in* δ_P *whenever full assignment is infeasible.*

(iii) ***Task values.*** *Fix* δ_A *and* δ_P*. If* F_H *first-order stochastically dominates* F_L*, then* $P^{N,K}(\delta_A, \delta_P, F_H) \geq P^{N,K}(\delta_A, \delta_P, F_L)$*.*

Greater worker patience makes favorable future assignments more effective compensation for unpleasant work today. Any assignment process that is feasible at a lower value of δ_A remains feasible after all promises are scaled down appropriately, so the firm can retain the same assignments. The firm can strictly improve whenever some tasks remain unassigned. As δ_A approaches zero, future compensation loses its force and the firm can assign only positive-value tasks in the limit. As δ_A approaches one, future assignments can support an arbitrarily long initial period of full assignment, so the firm's payoff approaches K. This limit does not require permanent full assignment to be feasible at any fixed δ_A: when $\mathbb{E}[V] \leq 0$, permanent full assignment is infeasible for every $\delta_A < 1$.

The effect of firm patience depends on the normalization of its objective. Let M_t be the number of tasks assigned at date t, and let T_{δ_P} be an independent random date with $\Pr(T_{\delta_P} = t) = (1 - \delta_P)\delta_P^t$. Under an optimal contract,

$$\frac{P^{N,K}(\delta_A, \delta_P, F)}{K} = \mathbb{E}\left[\frac{M_{T_{\delta_P}}}{K}\right].$$

Thus $P^{N,K}/K$ is the optimal discounted assignment fraction evaluated at a geometrically distributed date. A larger δ_P shifts weight toward later periods, when the firm may be fulfilling promises generated by earlier unpleasant assignments, and therefore lowers this normalized payoff. The unnormalized quantity $P^{N,K}/(1 - \delta_P)$ instead gives greater weight to every future assignment and strictly increases. Because a static contract delivers the same expected number of assignments in every period, its normalized payoff is independent of δ_P; the payoff gain from dynamic contracting therefore falls with firm patience.

For task distributions, couple the two economies by task ranks. Under the better distribution, the firm can reproduce the old matching and promise process while every assigned task is weakly more attractive to its recipient. The old assignment process therefore remains feasible, and the firm may be able to assign additional tasks. The proof of Proposition 9 is in Appendix A.4.

Long-run task assignment. The tasks completed over time are a central feature of how the employment relationship evolves. Two definitions are needed to characterize their long-run pattern. First, let $m_t(\mathbf{v})$ denote the matching the contract prescribes at date t if the realized task vector is $\mathbf{v}$, viewed from the start of the period before tasks are drawn. Second, I define a family of stationary benchmark rules. Fix a vector $\lambda \in \mathbb{R}_+^N$ specifying a marginal cost of promised utility for each worker, and suppose the firm faced these costs permanently. Its assignment problem would then be the same in every period: choose the matching that maximizes current completions, net of the cost—evaluated at λ—of the promises required to make each assignment acceptable. Let $\mathcal{A}_\lambda(\mathbf{v})$ be the set of matchings that solve this problem when the task vector is $\mathbf{v}$, after the continuation promise has been chosen optimally subject to (PC_{ik}). The exact objective is specified in Appendix A.5. The theorem asks whether the contract's date-t rule eventually chooses from one such fixed set of optimal matchings.

Almost all the conclusions below apply to every optimal Markov contract. The conclusion in part (iii) applies to *costate-assortative* matchings—those for which $\lambda_i < \lambda_j$ implies $Y_i \leq Y_j$. The optimal contract described in Theorem 2 is always costate-assortative.

Theorem 3 (Long-run task assignment when static full assignment fails)**.** *Suppose $K \leq N$, $\mathbb{E}[V] < \ell(\underline{v})$, and $u_0 = 0$. Every optimal Markov contract assigns every positive-valued task after almost every task draw at every reached history. The following long-run conclusions hold almost surely.*

(i) If $q > 1$, negative-valued tasks are assigned only finitely often.

(ii) If $q = 1$, positive- and negative-valued tasks are assigned with positive lower frequency.[4] *Costates can be chosen consistently so that $\lambda_t \to \lambda_\infty$ for a finite random vector λ_∞. For an independent task draw $\mathbf{v}$, the probability that the date-t rule selects a matching outside $\mathcal{A}_{\lambda_\infty}(\mathbf{v})$ converges to zero.*

(iii) If $q < 1$, positive- and negative-valued tasks are assigned with positive lower frequency. There exists an optimal costate-assortative contract for which every costate becomes positive in finite time, and with positive lower frequency, the recipient of the worst task moves from the lowest costate group to the highest in the next period.

The intuition is clearest if every worker's costate is strictly positive. The costate equation $\lambda_{t+1} = q\lambda_t + \mu_t/\delta_P$, where $\mu_t \geq 0$ and a multiplier can be positive only when a participation constraint binds, then gives a simple

[4] For a sequence of events $E_0, E_1, \ldots$, I say that E_t occurs with *positive lower frequency* if $\liminf_{T\to\infty} \frac{1}{T}\sum_{t=0}^{T-1} \mathbf{1}\{E_t\} > 0$ almost surely. Thus, along almost every path, the proportion of dates on which E_t occurs remains bounded away from zero over every sufficiently long horizon.

account of the three cases. If $q > 1$, every costate grows at least geometrically. Assigning a negative task requires a continuation promise, so only increasingly mild negative tasks can justify their growing marginal cost. Because task values have a bounded density, such assignments occur only finitely often almost surely. If $q = 1$, costates are nondecreasing and bounded, and hence converge to a finite, generally history-dependent limit λ_∞. The assignment rule then approaches the stationary benchmark associated with this realized limit. Because the limiting costate is finite, a fixed neighborhood of mildly negative task vectors remains worthwhile and recurs with positive frequency. If $q < 1$, runs of positive tasks have zero participation floors and shrink every costate by the factor q. These runs repeatedly return the contract to states in which promises are cheap enough that a positive-probability set of all-negative task vectors is assigned. Binding participation constraints then push costates upward. Under the costate-assortative selection, at the recurring dates identified in part (iii), the worker who takes the worst task moves from the lowest costate group to the highest in the next period.

The formal proof cannot assume that all costates eventually become positive. It begins at $\lambda_0 = 0$, and the value function can have kinks at which supporting costates and participation multipliers are not unique. Appendix A.5 first constructs finite costates and multipliers that satisfy the costate equation along any given optimal contract without changing its assignments or continuation promises. It then shows that at least K costates become positive when $q > 1$ and, for the selected costate-assortative contract, that every costate becomes positive when $q < 1$; under equal patience, the argument allows zero costates to persist. Tied costates also make recipient ordering nonunique, while the task-level conclusions must hold for every optimal Markov contract. Finally, the proof turns the preceding intuition into almost-sure statements. When $q > 1$, it shows that the probabilities of future negative assignments have a finite sum. When $q = 1$, it proves that costates are bounded and that the matching rule approaches its limiting problem. When $q < 1$, it shows that the process returns often enough to states with small costates that negative assignments occur with positive lower frequency.

A corollary of Theorem 3 is a characterization of the long-run fraction of all available tasks completed during the first T periods. If $q > 1$, this fraction converges almost surely to $\Pr(V > 0)$. If $q \leq 1$, its lower limit is almost surely strictly greater than $\Pr(V > 0)$. This observation helps resolve the surprising comparative static on firm patience in Proposition 9. Holding N, K, δ_A, and F fixed, raising δ_P from below δ_A to δ_A strictly lowers the firm's normalized discounted payoff but raises the long-run fraction of tasks completed. The normalized payoff falls because the more patient firm places greater weight on later periods, when it is fulfilling promises created by earlier unpleasant assignments; those same promises support the recurring completion of negative-valued tasks over a long horizon.

5 Discussion

This section first relates the model to assignment-based incentive mechanisms used in practice. Although the formal state variable is promised utility, organizations implement claims to favorable future assignments through queue positions, scheduling points, transfer points, and rotation rules; I discuss examples from gig economy platforms and civil service job allocation. I then consider two related problems: how outstanding promises interact with hiring and dismissal, and how the analysis changes when workers value the same task differently.

Gig economy platforms. Airport queues for rideshare drivers share several key features of the model: trips differ in desirability, fares do not price every relevant difference, drivers can decline offers, and platforms can condition future queue priority on current acceptance. A short trip from the airport illustrates the incentive problem. Although the fare reflects distance and time, the driver gives up a long queue wait for a small fare. Castro et al. (2021) study exactly this problem, showing that under first-in-first-out dispatch, heterogeneity in trip earnings leads drivers to cherry-pick and decline short trips, and that queue-position mechanisms can realign incentives. The platforms' responses take the form the theory predicts. Uber's "Priority Queue Access" holds a driver's place in the queue after a short airport trip;[5] Lyft's "short ride bump" upgrades the queue position of a driver who returns after accepting a short ride;[6] DiDi places drivers who complete short airport trips into a priority queue on their return.[7] Accepting a bad draw today earns a claim on a better position tomorrow.

Amazon Flex operates a more explicit point system. Drivers earn points for completed deliveries and delivery blocks, and higher reward levels unlock "Preferred Scheduling": reserved offers at the driver's preferred stations, days, and times.[8] The instrument is the one the theory identifies—accumulated points convert past work into priority over future assignments—although Flex points reward the volume of past work rather than its difficulty.

Civil service job allocation. Where pay is set by legislated scales, pay differentials may exist—remote area allowances, hardship pay, sea pay—but they are often fixed by schedule rather than adjusted until postings fill. As a result, hard-to-staff postings persist. Assignment-based instruments carry much of the remaining incentive load. Queensland's Department of Education rates schools by remoteness, and teachers in remote

[5] https://help.uber.com/en-AU/driving-and-delivering/article/queue-access?nodeId=1275528f-2ba7-4ca4-982f-6aca4e3fad99

[6] https://help.lyft.com/hc/en-us/all/articles/115013080368-Pro-Tips-for-Giving-Airport-Rides

[7] https://web.didiglobal.com/au/help-center/driving-with-didi-at-the-airport/

[8] https://flex.amazon.com/amazonflexrewards

schools accrue *transfer points* several times faster than teachers in metropolitan schools; accumulated points determine the priority of transfer requests to preferred locations.[9] The U.S. Foreign Service requires service at high-hardship posts of officers competing for entry into the Senior Foreign Service,[10] and the U.S. Navy's "Sea Shore Flow" policy prescribes how demanding sea tours are offset by shore duty.[11] In each case, a difficult posting earns a claim to a better one—a claim the pay scale itself does not provide.

Workforce adjustment. The model holds the workforce size and employment arrangement fixed. In practice, firms can hire additional workers, dismiss existing workers, and substitute temporary for permanent employment. Because the wage bill is constant when the workforce is fixed, the firm's objective omits the flow cost Nw; the workforce comparisons in Proposition 7 therefore measure the gross assignment benefit of additional workers. Starting from zero promises, the firm's normalized net payoff with N permanent workers is $\Phi^{N,K}(0) - Nw$. Define the marginal gross benefit of worker $N+1$ by $\Delta_N^K := \Phi^{N+1,K}(0) - \Phi^{N,K}(0)$. The option to leave the new worker idle gives $\Delta_N^K \geq 0$. In the absence of adjustment costs, the firm strictly prefers to hire that worker when $\Delta_N^K > w$, strictly prefers not to hire when $\Delta_N^K < w$, and is indifferent at equality. Proposition 7 gives $\Delta_N^1 = 0$ when one task arrives each period and $\Delta_N^K = 0$ once the existing workforce supports full assignment. With several simultaneous tasks, Δ_N^K can be strictly positive because the additional worker provides spare capacity while another worker is owed favorable treatment. The firm may therefore optimally employ more workers than there are tasks when this spare-capacity benefit exceeds the wage.

The same calculation describes firing when the worker has no outstanding promise. At a state u with $u_i = 0$ and $N - 1 \geq K$, the marginal gross value of retaining worker i is $\Delta_i^K(u) := \Phi^{N,K}(u) - \Phi^{N-1,K}(u_{-i})$; absent dismissal costs, the firm dismisses her when $\Delta_i^K(u) < w$. A worker with positive promised utility has already earned a claim on favorable future assignments, so dismissing her and evaluating $\Phi^{N-1,K}(u_{-i})$ would erase a contractual obligation. That possibility would weaken the incentive to accept difficult tasks earlier in the relationship. A full analysis of workforce adjustment therefore requires assumptions about severance, the vesting of assignment promises, the firm's commitment to retain workers, and the costs of hiring and dismissal. Endogenizing these choices is a natural direction for future research.

Workforce heterogeneity. The model also assumes that every worker assigns the same value to a given task. Much of the recursive approach extends when the firm instead observes a value matrix (V_{ik}), where V_{ik} is worker i's value for task k. Promised utilities remain sufficient state variables, constraint (PC_{ik}) for match

[9] https://teach.qld.gov.au/studyteaching/Documents/remote-teaching-booklet.pdf

[10] https://afsa.org/serving-hardship-posts

[11] https://www.mynavyhr.navy.mil/Career-Management/Community-Management/Enlisted/Sea-Shore-Flow/

(i, k) uses the floor $\ell(V_{ik})$, and the costate equation continues to describe how the marginal costs of promises evolve. A particularly close extension assumes that the joint distribution of the value matrix is invariant to relabelling workers: for every permutation π, the matrices $(V_{\pi(i)k})$ and (V_{ik}) have the same distribution. Under this exchangeability condition, the feasible promise set and value function remain symmetric, so concavity again implies Schur-concavity and the firm retains its preference for spreading promises evenly.

The realized allocation has a precise matching interpretation. Fix a costate vector λ and a candidate continuation promise z. Conditional on z, the firm chooses a maximum-weight matching among edges satisfying $z_i \geq \ell(V_{ik})$, with edge weight $a_{ik} := (1-\delta_P) + (1-\delta_A)\lambda_i V_{ik}$. It then chooses z jointly with the matching, taking account of the continuation contribution $\delta_P \Phi(z) + \delta_A \lambda \cdot z$. Because this continuation contribution links the workers' promises, optimizing out z need not leave an additive edge-weight problem. Realized comparative advantage and the marginal costs of promises therefore jointly determine the assignment. General worker-specific values need not produce an assortative matching because workers may rank the same tasks differently. A common task ranking together with compatible increasing differences in the adjusted edge weights and participation floors could restore assortativity. With permanent worker types, exchangeability and Schur-concavity continue to apply within each type. The recursive structure survives more generally, while the common cutoff, priority ranking, and long-run task-level conclusions require further analysis. Establishing those extensions is a useful direction for future research.

6 Conclusion

The central lesson of the model is that task assignment can serve as compensation. When pay cannot adjust task by task, a firm can reward difficult work through the tasks a worker is assigned later. Promised utility acts as a balance in an internal account. Accepting an unpleasant task credits the account by increasing the future utility the firm owes the worker. Receiving a desirable task, or being protected from an unpleasant one, draws down the balance by delivering some of that utility. The assignment rule determines when the balance is built and how it is repaid.

This analysis helps explain how history and patience shape the employment relationship. Each assignment changes what the firm owes the worker in the future, and a larger balance makes the firm more selective about imposing further difficult work. Relative patience determines how these employment relationships evolve. Promised utility may accumulate into seniority, settle into a stable assignment rule, or be repeatedly built through difficult work and repaid through favorable treatment. The same incentive problem can therefore produce protected status, stable priorities, or rotation between harder and easier assignments.

With several workers, these obligations interact through the matching problem. Desirable tasks help the firm compensate workers who are owed more, while unpleasant tasks can be directed toward workers who are owed less. Pooling task streams gives the firm more ways to match assignments to these obligations, and a spare worker can preserve capacity while another worker is owed favorable treatment. Together, flexible matching and slack capacity help the firm honor what it owes while continuing to assign tasks across the workforce.

This model overlooks several features that matter in practice. Task costs may be privately observed, firms may be unable to commit fully to promised priority, and workforce turnover may end relationships before outstanding obligations are repaid. Extending the analysis along these dimensions would show how information, commitment, and turnover shape assignment-based incentives. Administrative data from queueing, transfer, and rotation systems used by firms could help guide and test those extensions. The power of the present model lies in showing that substantial dynamics arise from two simple constraints: pay cannot adjust task by task, and workers can refuse assignments. These constraints alone make past assignments matter for future treatment. The central conclusion—that assignment rules record and repay the obligations created by difficult work—is therefore likely to extend to richer models of organizations.

A Omitted Proofs

A.1 Recursive Formulation and the Elimination of Randomization

This subsection establishes recursive feasibility and shows that lotteries conditional on the task realization do not raise the principal's value. A promise set is *self-generating* if every promise in it can be delivered by a one-period rule whose continuation promises remain in the set.

Let $\partial^+\mathcal{U}^{N,K} := \{u \in \mathcal{U}^{N,K} : \sum_i u_i = \overline{V}_N\}$ denote the upper boundary face of the promise polytope. The single-agent model follows by setting $N = K = 1$.

Lemma 1 (The promise set is self-generating)**.** *For every $u \in \mathcal{U}^{N,K}$, there is a possibly randomized one-period rule that delivers each worker at least u_i and always chooses the next promise in $\mathcal{U}^{N,K}$.*

Proof. For each permutation π, let P_π be the permutation matrix satisfying $(P_\pi x)_{\pi(r)} = x_r$, and set $b^\pi = P_\pi \bar{v}$. A standard majorization theorem gives, for every $u \preceq_w \bar{v}$, a vector $\tilde{u} \geq u$ that is majorized by $\bar{v}$ and has the same total as $\bar{v}$ (see Marshall, Olkin, & Arnold, 2011). Every $x \in \partial^+\mathcal{U}^{N,K}$ is majorized by $\bar{v}$ in the usual strong sense. By the Hardy–Littlewood–Pólya theorem, $x = A\bar{v}$ for some doubly stochastic matrix A. By the Birkhoff–von Neumann theorem, there are coefficients $\alpha_\pi \geq 0$, with $\sum_\pi \alpha_\pi = 1$, such that $A = \sum_\pi \alpha_\pi P_\pi$. Hence $x = \sum_\pi \alpha_\pi b^\pi$. Conversely, each b^π belongs to $\partial^+\mathcal{U}^{N,K}$, and this face is convex.

Therefore $\partial^{+}\mathcal{U}^{N,K} = \mathrm{co}\{b^{\pi}\}$. Applying this representation to $\tilde{u}$ gives coefficients α_{π} such that $\tilde{u} = \sum_{\pi} \alpha_{\pi} b^{\pi}$.

Draw π publicly with probabilities α_{π} before the task vector, assign the nonnegative tasks in descending order to workers $\pi(1), \pi(2), \ldots$, and set the next promise to b^{π}. Conditional on π, worker $\pi(r)$'s expected current utility is $\bar{v}_r$, so the expected current-utility vector is b^{π}. Writing Y for the vector of current task utilities, $\mathbb{E}[(1-\delta_A)Y + \delta_A b^{\pi} \mid \pi] = b^{\pi}$. Every assigned task is nonnegative, so (PC_{ik}) holds. The lottery therefore delivers $\tilde{u} \geq u$ and keeps every continuation in $\mathcal{U}^{N,K}$, proving feasibility at u. □

Let $\mathcal{M}^{N,K}$ denote the finite set of deterministic matchings and define $\mathcal{Z}(m,v) := \{z \in \mathcal{U}^{N,K} : z_i \geq \ell(v_k) m_{ik} \text{ for all } i,k\}$. To establish existence, temporarily enlarge the Bellman problem in (MOPT): after observing v, the principal may publicly randomize over pairs $(m,z) \in \mathcal{M}^{N,K} \times \mathcal{Z}(m,v)$. This *relaxed problem* has the same objective and promise-keeping inequalities as the original problem; it differs only by permitting this one-period lottery.

Equivalently, a relaxed decision is a joint probability law for the task vector, matching, and continuation promise. Under this law, the task vector must have marginal distribution F^K; the law must be supported on triples (v,m,z) satisfying $z \in \mathcal{Z}(m,v)$; and each worker's expected delivered utility must be at least u_i.

Lemma 2 (Existence and regularity for the relaxed problem)**.** *The relaxed Bellman operator has a unique bounded fixed point $\Phi^{N,K}$. This fixed point is concave and continuous on $\mathcal{U}^{N,K}$, and the relaxed maximum is attained at every promise vector.*

Proof. Let $\Gamma(u)$ be the set of relaxed decisions feasible at u. The space $\mathcal{V}^K \times \mathcal{M}^{N,K} \times \mathcal{U}^{N,K}$ is compact, and the set of triples satisfying $z \in \mathcal{Z}(m,v)$ is closed because ℓ is continuous. Consequently, Γ has nonempty compact values and a closed graph; nonemptiness follows from Lemma 1. For every bounded upper-semicontinuous continuation value, the objective is upper semicontinuous in the joint law. The upper-semicontinuous maximum theorem therefore implies that the Bellman operator attains its maximum and preserves upper semicontinuity.

The operator is a δ_P-contraction in the sup norm. Since it maps the complete space of bounded upper-semicontinuous functions into itself, iteration from zero converges uniformly to a bounded upper-semicontinuous fixed point $\Phi^{N,K}$. The contraction inequality makes this the unique bounded fixed point. Publicly randomizing between any two feasible laws proves that $\Phi^{N,K}$ is concave. Concavity gives lower semicontinuity in the relative interior, meaning the interior within the smallest affine space containing the promise set. Now fix a boundary point u and a sequence $u_n \to u$. Because $\mathcal{U}^{N,K}$ is a polytope described by finitely many linear inequalities, there is a constant $c(u) > 0$ such that the feasible ray from u through any $u_n \neq u$ can be extended to a point $b_n \in \mathcal{U}^{N,K}$ satisfying $\|b_n - u\| \geq c(u)$. Writing $u_n = (1-\varepsilon_n)u + \varepsilon_n b_n$, we have $\varepsilon_n \to 0$, so concavity and boundedness give $\liminf_n \Phi^{N,K}(u_n) \geq \Phi^{N,K}(u)$. Thus $\Phi^{N,K}$ is continuous. $\square$

Lemma 3 (One-period randomization can be eliminated)**.** *The value of the relaxed Bellman problem is attained by deterministic matching and continuation-promise rules. Consequently, the deterministic Bellman problem* (MOPT) *has the same value. The deterministic rule can also preserve the expectations of any specified finite collection of bounded measurable statistics of the task vector and matching. For $N = K = 1$, this gives the scalar specialization of* (MOPT)*.*

Proof. The measurable maximum theorem and the disintegration theorem allow the optimal relaxed rule to be chosen so that its conditional matching probabilities $q_m(v, u)$ and conditional continuation distributions are jointly measurable in (v, u). Conditional on (v, u, m), replace the lottery over $\mathcal{Z}(m, v)$ by its mean $z_m(v, u)$. Since $\mathcal{Z}(m, v)$ is convex, $z_m(v, u) \in \mathcal{Z}(m, v)$ whenever $q_m(v, u) > 0$. This preserves feasibility, and concavity of $\Phi^{N,K}$ weakly raises the value; optimality makes the increase zero.

Extend z_m arbitrarily wherever $q_m(v, u) = 0$. Since F has a density, F^K is atomless. Thus, the finite-action purification theorem of Dvoretzky, Wald, and Wolfowitz (1951) allows us to replace the measurable lottery over finitely many matchings by a measurable deterministic matching while preserving any specified finite collection of expectations. Apply it to the vector consisting of the principal's payoff conditional on (v, u, m), the N delivered-utility components, the specified matching statistics, and $a(v, u, m) := \mathbf{1}\{q_m(v, u) = 0\}$. It gives a deterministic matching $m^*(v, u)$ preserving all these expectations.[12] The randomized expectation of a is zero, so $q_{m^*(v,u)}(v, u) > 0$ almost surely. Pair m^* with $u^*(v, u) := z_{m^*(v,u)}(v, u)$. On the null set where $q_{m^*(v,u)}(v, u) = 0$, use the empty matching and continuation promise zero. This changes no expectation and ensures (PC$_{ik}$) for every task vector. $\square$

A.2 Value Function Properties

This section proves Proposition 6; setting $N = K = 1$ gives Proposition 1.

Basic value function properties.

[12]The deterministic choice can be made universally measurable in u. Represent each matching by the corresponding unit vector in $\mathbb{R}^{|\mathcal{M}^{N,K}|}$. A deterministic matching rule is then an element of $L^1(F^K; \mathbb{R}^{|\mathcal{M}^{N,K}|})$ that takes values in these unit vectors. For each u, the set of rules that preserve the specified expectations and select only matchings with positive probability under the original lottery is nonempty and closed, and these sets have a Borel graph. The Jankov–von Neumann selection theorem therefore gives a universally measurable L^1-valued choice. It has a jointly measurable representative after completing the state and task sigma-fields under the induced laws.

Proof of Proposition 6(i). Lemma 2 gives a unique bounded relaxed fixed point $\Phi^{N,K}$, which is concave and continuous. Lemma 3 shows that this function also satisfies the deterministic Bellman equation. The deterministic Bellman operator obeys the same δ_P-contraction bound in the sup norm, so $\Phi^{N,K}$ is its unique bounded fixed point.

If $u \le \widetilde{u}$, every policy feasible at $\widetilde{u}$ is feasible at u, so $\Phi^{N,K}(u) \ge \Phi^{N,K}(\widetilde{u})$. Relabeling workers maps every feasible policy at u to a feasible policy at $P_\pi u$ with the same payoff, so $\Phi^{N,K}$ is symmetric. If x majorizes y, the convex-hull characterization of majorization gives weights $\alpha_\pi \ge 0$, with $\sum_\pi \alpha_\pi = 1$, such that $y = \sum_\pi \alpha_\pi P_\pi x$ (see Marshall, Olkin, & Arnold, 2011). Concavity and symmetry then give $\Phi^{N,K}(y) \ge \sum_\pi \alpha_\pi \Phi^{N,K}(P_\pi x) = \Phi^{N,K}(x)$. Thus $\Phi^{N,K}$ is Schur-concave. □

Value function geometry. The upper face of the promise set exhausts the economy's capacity to deliver worker utility. Relaxing a positive coordinate from this face allows a small increase in the principal's payoff, so the marginal cost of promised utility becomes unbounded there. The next result formalizes this observation and will also show that the plateau ends before the upper face.

Lemma 4 (No finite costate on the upper face)**.** *If $u \in \mathcal{U}^{N,K}$ satisfies $\sum_i u_i = \overline{V}_N$, then no finite vector supports $-\Phi^{N,K}$ at u. In particular, $\Lambda^{N,K}(u) = \varnothing$. Thus, a path on which a finite supporting costate is selected at every date cannot reach the upper face.*

Proof. Summing promise keeping over agents, and writing Y and Z for total current utility and the total continuation promise, gives $\overline{V}_N \le \mathbb{E}[(1-\delta_A)Y + \delta_A Z] \le (1-\delta_A)\overline{V}_N + \delta_A \overline{V}_N$. Here $Y \le \sum_k [V_k]_+$ pointwise, $\mathbb{E}[\sum_k [V_k]_+] = \overline{V}_N$, and $Z \le \overline{V}_N$. Equality throughout forces the current allocation to assign every positive task and no negative task, and the next promise to remain on the upper face almost surely. These equality conditions recur at every continuation state, so every feasible contract from the upper face has payoff at most $K\,\Pr(V > 0)$. The representation in Lemma 1 supplies a public lottery over worker orders whose expected current-utility vector is u. Assigning positive tasks in that order and retaining continuation promise u delivers $(1-\delta_A)u + \delta_A u = u$ and attains this bound. Hence $\Phi^{N,K}(u) = K\,\Pr(V > 0)$.

Suppose, toward a contradiction, that a finite λ supports $-\Phi^{N,K}$ at u. Choose j with $u_j > 0$ and then $\varepsilon > 0$ such that $\ell(-\varepsilon) < u_j$, $(-\varepsilon, -\varepsilon/2) \subset \mathcal{V}$, and $1 - \delta_P > \max\{\lambda_j, 0\}(1-\delta_A)\varepsilon$. Let E_ε be the event that all task values lie in this interval, which has probability $p_\varepsilon > 0$ by the positive-density assumption. The preceding rule assigns nothing on this event. Assign task 1 to worker j instead and retain continuation u. Constraint (PC_{ik}) holds because $\ell(V_1) \le \ell(-\varepsilon) < u_j$ on E_ε. Relative to the upper-face rule just constructed, this modification changes the expected delivered-utility vector from u to $u - d_\varepsilon e_j$, where

$d_\varepsilon := (1-\delta_A)\mathbb{E}[-V_1 \mathbf{1}_{E_\varepsilon}] \le (1-\delta_A)\varepsilon p_\varepsilon < u_j$, and leaves the continuation value equal to $\Phi^{N,K}(u)$. Downward closure makes $u - d_\varepsilon e_j$ feasible, while the extra assignment gives

$$\Phi^{N,K}(u - d_\varepsilon e_j) - \Phi^{N,K}(u) \ge (1-\delta_P)p_\varepsilon.$$

The supporting inequality bounds the left side by $\lambda_j d_\varepsilon$. A negative λ_j gives an immediate contradiction. If $\lambda_j \ge 0$, the bound on d_ε gives $1-\delta_P \le \lambda_j(1-\delta_A)\varepsilon$, contradicting the choice of ε. □

Lemma 5 (Plateau)**.** *Let $\mathcal{P}^{N,K} := \{u \in \mathcal{U}^{N,K} : \Phi^{N,K}(u) = \Phi^{N,K}(0)\}$. Then $\mathcal{P}^{N,K}$ is nonempty, compact, convex, symmetric, and downward closed: if $u \in \mathcal{P}^{N,K}$ and $0 \le u' \le u$, then $u' \in \mathcal{P}^{N,K}$. It is a strict subset of $\mathcal{U}^{N,K}$ and contains a nonzero promise vector. For $N = K = 1$, $\mathcal{P}^{1,1} = [0,\bar{u}]$ for some $\bar{u} \in (0, V^+)$.*

Proof. Because every promise vector is nonnegative and $\Phi^{N,K}$ is nonincreasing, $\Phi^{N,K}(0)$ is its maximal value. Thus $0 \in \mathcal{P}^{N,K}$. Compactness follows from continuity and symmetry from symmetry of $\Phi^{N,K}$. The set on which a concave function attains its maximum is convex, and monotonicity makes this set downward closed.

Let $(m(v), z(v))$ be an optimal one-period rule at $u = 0$, and let g be its expected delivered-utility vector, with $g_i := \mathbb{E}[(1-\delta_A)\sum_k V_k m_{ik}(\mathbf{V}) + \delta_A z_i(\mathbf{V})]$. Constraint (PC$_{ik}$) makes every coordinate of the realized delivered-utility vector nonnegative: the continuation lower bound offsets the current disutility of an assigned negative task, while a positive task or idleness also delivers nonnegative utility. The rule assigns every positive task F^K-almost surely. If it left a positive task unassigned on a set of positive probability, at most $K-1$ workers would be occupied on that set. Since $N \ge K$, assigning the task to an idle worker and retaining the continuation promise would preserve (PC$_{ik}$), weakly increase delivered utility, and strictly raise the principal's expected payoff. Positive tasks occur with positive probability, so $\sum_i g_i > 0$, and hence $g_i > 0$ for some worker i.

The same rule is feasible at $g_i e_i$ and earns $\Phi^{N,K}(0)$, so $g_i e_i \in \mathcal{P}^{N,K}$. Symmetry gives $g_i e_j \in \mathcal{P}^{N,K}$ for every worker j, and convexity then gives $(g_i/N)\mathbf{1} \in \mathcal{P}^{N,K}$. Thus the plateau contains a strictly positive promise vector.

Suppose $u \in \partial^+\mathcal{U}^{N,K} \cap \mathcal{P}^{N,K}$. Then u minimizes $-\Phi^{N,K}$ on $\mathcal{U}^{N,K}$, so $0 \in \partial(-\Phi^{N,K})(u)$, and therefore $0 \in \Lambda^{N,K}(u)$. This contradicts Lemma 4. Hence $\partial^+\mathcal{U}^{N,K} \cap \mathcal{P}^{N,K} = \varnothing$, which makes $\mathcal{P}^{N,K}$ a strict subset of $\mathcal{U}^{N,K}$.

When $N = K = 1$, the state space is $[0, V^+]$. A nonempty compact convex downward closed strict subset of $[0, V^+]$ containing a nonzero point is exactly $[0, \bar{u}]$ with $\bar{u} \in (0, V^+)$. □

The remaining lemmas study $\Lambda^{-1}(\lambda)$, the set of promise vectors at which λ supports $-\Phi^{N,K}$. We introduce the one-period Lagrangian notation used to compare optimal contracts at any two promises in this set.

The comparison requires some notation. Fix a costate vector λ and a task vector v. For a matching $m \in \mathcal{M}^{N,K}$, write $Y_i(m,v) := \sum_k v_k m_{ik}$ for agent i's current utility and $\ell_i(m,v) := \sum_k \ell(v_k) m_{ik}$ for the smallest continuation promise that makes the assigned task acceptable to agent i. For a feasible matching and continuation promise (m,z), write $G(m,z;v) := (1-\delta_A)Y(m,v) + \delta_A z$ for the utility vector delivered toward the current promises. Define $A_\lambda(z) := \delta_P \Phi^{N,K}(z) + \delta_A \lambda \cdot z$, the part of the current Lagrangian determined by the continuation promise, and set $\Gamma_\lambda(\ell) := \max\{A_\lambda(z) : z \in \mathcal{U}^{N,K},\ z \geq \ell\}$, with value $-\infty$ when no such z exists. Thus, $\Gamma_\lambda(\ell)$ is the largest continuation contribution among promises that make the current assignments acceptable. After optimizing over the continuation promise, the current λ-Lagrangian is

$$W_\lambda(m;v) := (1-\delta_P)\sum_{i,k} m_{ik} + (1-\delta_A)\sum_i \lambda_i Y_i(m,v) + \Gamma_\lambda(\ell(m,v)).$$

This expression ranks matchings by the principal's current assignment payoff and by the utility delivered toward the workers' promises, valued at the marginal promise costs λ, after choosing the best feasible continuation promise. A matching is λ-optimal at v if it maximizes $W_\lambda(\cdot;v)$ over $\mathcal{M}^{N,K}$. A feasible pair (m,z) is λ-optimal at v if m is λ-optimal and z attains $\Gamma_\lambda(\ell(m,v))$. A λ-optimal one-period policy is a measurable rule $v \mapsto (m(v), z(v))$ such that $m(v)$ is λ-optimal and $z(v)$ attains $\Gamma_\lambda(\ell(m(v),v))$ after every v. A matching is λ-assortative if $Y_i(m,v) \leq Y_j(m,v)$ whenever $\lambda_i < \lambda_j$. Recall that an equal-costate group contains the workers with the same coordinate of λ, and $S_\lambda x$ records the sum of x_i within each such group. Relabeling workers within one group does not change $S_\lambda Y(m,v)$.

Standard Lagrangian optimality for promise keeping implies that, for any $u \in \mathcal{U}^{N,K}$, any Bellman-optimal one-period rule $(m(v), z(v))$ at u, and every $\lambda \in \Lambda^{N,K}(u)$, the selected matching and continuation promise are λ-optimal for F^K-almost every v. A measurable λ-optimal selection can replace the rule on the exceptional null set, yielding a λ-optimal one-period policy. If $g := \mathbb{E}[G(m(\mathbf{V}), z(\mathbf{V}); \mathbf{V})]$, complementary slackness gives $\lambda_i(g_i - u_i) = 0$ for every worker i. In particular, every promise-keeping constraint binds when $\lambda \in \mathbb{R}^N_{++}$.

Lemma 6 (λ-assortative current assignment)**.** *Fix $\lambda \in \mathbb{R}^N_+$. Any λ-optimal one-period policy can be modified so that, after every task vector, its matching assigns every positive task, is λ-assortative, and, if it assigns negative tasks, assigns the least negative ones down to a cutoff. The modified policy remains measurable and λ-optimal, preserves the principal's Bellman payoff, and weakly increases G coordinatewise.*

Proof. Fix a task vector v, and write (m,z) for the matching and continuation promise selected by the policy.

Every positive task is assigned: if one were unassigned, at most $K - 1$ workers would be occupied, so $N \geq K$ leaves an idle worker who could receive it. Retaining z would preserve (PC_{ik}) and strictly raise the realized Lagrangian, contradicting optimality. If a negative task x is assigned while a task $y \in (x, 0]$ is unassigned, replace x by y for the same worker and retain z. This leaves the principal's Bellman payoff after v unchanged, relaxes that worker's participation constraint, and weakly raises her delivered utility. The modified matching and continuation promise are therefore also λ-optimal. Repeating this replacement finitely many times gives a λ-optimal matching and continuation promise that assign the least negative tasks down to a cutoff.

Now sort the selected tasks across workers. Append one dummy task with value and participation floor zero for each idle worker, so the same argument covers agents who are idle. Consider workers i, j with $\lambda_i < \lambda_j$, where i receives value y and j receives value $x < y$. Let $\widetilde{m}$ swap these values, and set $a := \ell(x)$ and $b := \ell(y)$, so $a \geq b$. Feasibility gives $z_i \geq b$ and $z_j \geq a$.

We first show that $z_i \leq z_j$. Otherwise, swapping the two task–continuation packages preserves feasibility and, by symmetry, leaves $\Phi^{N,K}$ unchanged. It raises the realized λ-Lagrangian by $(\lambda_j - \lambda_i)[(1 - \delta_A)(y - x) + \delta_A(z_i - z_j)] > 0$, contradicting optimality.

Put $d := (a - z_i)_+$. Since $z_i \geq b$, $z_j \geq a$, and $z_i \leq z_j$, we have $d \leq a - b$ and $d \leq z_j - z_i$. Hence $\widehat{z} := z + d(e_i - e_j)$ lies on the segment between z and its (i, j)-coordinate swap, so symmetry and convexity of $\mathcal{U}^{N,K}$ give $\widehat{z} \in \mathcal{U}^{N,K}$. Moreover, $\widehat{z}_i \geq a$ and $\widehat{z}_j \geq b$, so $\widehat{z}$ satisfies the lower bounds for $\widetilde{m}$.

The vector z majorizes $\widehat{z}$, so Schur-concavity gives $\Phi^{N,K}(\widehat{z}) \geq \Phi^{N,K}(z)$. The definition of ℓ also gives $d \leq a - b \leq ((1 - \delta_A)/\delta_A)(y - x)$. Since the swap leaves the number of assigned tasks unchanged and $\widehat{z}$ is feasible for $\widetilde{m}$,

$$0 \geq W_\lambda(\widetilde{m}; v) - W_\lambda(m; v) \geq \delta_P \left[\Phi^{N,K}(\widehat{z}) - \Phi^{N,K}(z)\right] + (\lambda_j - \lambda_i)\left[(1 - \delta_A)(y - x) - \delta_A d\right] \geq 0.$$

Thus equality holds throughout: $\widetilde{m}$ is λ-optimal, $\widehat{z}$ attains $\Gamma_\lambda(\ell(\widetilde{m}, v))$, and $\Phi^{N,K}(\widehat{z}) = \Phi^{N,K}(z)$. Hence $\widetilde{m}$ together with $\widehat{z}$ is λ-optimal. Because $\lambda_i < \lambda_j$, equality also gives $\delta_A d = (1 - \delta_A)(y - x)$, which means that the two affected coordinates of G are unchanged. The task count is unchanged, and the equality of continuation payoffs preserves the principal's Bellman payoff after v.

Repeat this exchange whenever the assignment contains an inversion. Each exchange reduces the finite number of inversions, so the process produces a λ-assortative matching after finitely many steps. Choosing the next inverted worker pair according to a fixed order makes the modified policy measurable in v. The task improvements weakly raise G, while the sorting exchanges leave it unchanged. Thus the modified policy remains λ-optimal, preserves the principal's Bellman payoff, and weakly raises G coordinatewise. If the original policy is Bellman-optimal at a promise, these properties preserve feasibility and optimality at that promise. Relabeling workers within an equal-costate group preserves the group sum $S_\lambda G$ and may redistribute its individual coordinates. $\square$

Lemma 7 (Uniqueness of group-level task assignment)**.** *Fix $\lambda \in \mathbb{R}^N_{++}$. For F^K-almost every task vector, every λ-optimal pair (m, z) can be rearranged into a λ-optimal λ-assortative pair (m^a, z^a) satisfying $S_\lambda G(m^a, z^a; v) = S_\lambda G(m, z; v)$. Moreover, all λ-optimal λ-assortative pairs assign the same multiset of task values to each equal-costate group, up to relabeling workers within that group.*

Proof. Because F has a density, task values are distinct and nonzero almost surely. Fix such a task vector. Every λ-optimal pair assigns all positive tasks by the idle-worker argument in the preceding proof. Moreover, if value x is assigned to worker i while $y > x$ is unassigned, replacing x by y for the same worker and retaining the continuation promise relaxes that worker's participation constraint and raises the realized Lagrangian by at least $(1 - \delta_A)\lambda_i(y - x) > 0$. Thus, conditional on assigning k negative tasks, every optimum assigns the k least negative values.

The sorting construction in Lemma 6 uses exchanges between workers with distinct costates, each of which preserves G coordinatewise. Relabeling workers within an equal-costate group preserves $S_\lambda G$, which gives the claimed rearrangement. For a fixed k, append one zero-valued dummy task for each idle worker and order the resulting list of N actual and dummy tasks by value. Assortativity assigns successive blocks of this ordered list to the equal-costate groups, ordered by costate. The strict ordering of the nonzero task values and the sizes of the groups therefore determine the multiset assigned to each group. It remains to show that k is unique.

On the event that there are M negative values, order them as $y_1 > \cdots > y_M$, and let $H_k(v; \lambda)$ be the largest value of $W_\lambda(m; v)$ among matchings whose assigned task set consists of every positive task and $y_1, \ldots, y_k$. Set $H_k(v; \lambda) = -\infty$ when no such matching is feasible. Fix $k > \ell$, condition on M, the positive order statistics, and every negative order statistic except y_k, and view H_k and H_ℓ as functions of y_k. The value H_ℓ is independent of y_k. Whenever H_k is finite, deleting tasks $y_{\ell+1}, \ldots, y_k$ shows that H_ℓ is also finite. If y_k rises from t to t' without crossing a neighboring order statistic and $H_k(t; \lambda)$ is finite, retain an optimizer at t, assign t' to the same worker i, and retain the continuation promise. The participation floor weakly falls, while the realized Lagrangian rises by $(1 - \delta_A)\lambda_i(t' - t) > 0$. Thus $H_k - H_\ell$ is strictly increasing wherever H_k is finite. The conditional distribution of y_k has a density on its admissible interval, so $\Pr(H_k = H_\ell > -\infty) = 0$. There are finitely many pairs $k > \ell$, and H_0 is finite because all positive tasks can be assigned with zero continuation promises. Hence the maximizing count is unique almost surely, which proves the group-level uniqueness claim. ◻

Lemma 8 (Costate for the continuation promise)**.** *Fix a matching $m \in \mathcal{M}^{N,K}$. For F^K-almost every task vector v, every $\lambda \in \mathbb{R}_+^N$ and every continuation promise z attaining $\Gamma_\lambda(\ell(m,v))$ admit a multiplier $\alpha \in \mathbb{R}_+^N$ on (PC$_{ik}$) and a next-period costate $\lambda^+ \in \Lambda^{N,K}(z)$ satisfying $\alpha_i(z_i - \ell_i(m,v)) = 0$ for every i and $\lambda^+ = q\lambda + \alpha/\delta_P$.*

Proof. Fix a matching m and write $e(v) := \ell(m,v)$. Suppose $e(v)$ is feasible. If $\sum_{i\in S} e_i(v) < \overline{V}_{|S|}$ for every nonempty set of workers S, then $e(v) + \varepsilon\mathbf{1}$ lies in the relative interior of $\mathcal{U}^{N,K}$ for all sufficiently small $\varepsilon > 0$. It also strictly exceeds every participation floor. Thus the relative Slater condition holds whenever all the subset constraints have positive slack at $e(v)$.

For a fixed nonempty S, the boundary constant $\overline{V}_{|S|}$ is positive. Conditional on the signs of the distinct, independent task values assigned to workers in S, the positive participation floors are continuously distributed. Their sum is either zero or has a density, so it equals $\overline{V}_{|S|}$ with probability zero. There are finitely many worker subsets, so for F^K-almost every v, the floor vector $e(v)$ satisfies relative Slater whenever it is feasible.

Fix such a v, any $\lambda \in \mathbb{R}_+^N$, and a continuation promise z attaining $\Gamma_\lambda(e(v))$. Standard Kuhn–Tucker conditions give $\alpha \in \mathbb{R}_+^N$, with $\alpha_i(z_i - e_i(v)) = 0$, such that z maximizes the continuation Lagrangian $\delta_P\Phi^{N,K}(\tilde{z}) + \delta_A\lambda\cdot\tilde{z} + \alpha\cdot(\tilde{z} - e(v))$ over $\tilde{z} \in \mathcal{U}^{N,K}$. Dropping the constant term and dividing by δ_P shows that $\lambda^+ := q\lambda + \alpha/\delta_P$ belongs to $\Lambda^{N,K}(z)$, which gives the stated costate equation. ◻

Lemma 9 (What a strictly positive costate determines)**.** *Fix $\lambda \in \mathbb{R}_{++}^N$. All promise vectors in $\Lambda^{-1}(\lambda)$ have the same total promise within each equal-costate group; equivalently, $S_\lambda\Lambda^{-1}(\lambda)$ is a singleton. In particular, if the coordinates of λ are all distinct, then $\Lambda^{-1}(\lambda)$ is a singleton. For $N = K = 1$, this implies Φ is strictly concave on $(\bar{u}, V^+]$.*

Proof. For every $\lambda \in \mathbb{R}_+^N$, the set $\Lambda^{-1}(\lambda)$ is the set of maximizers of $\Phi^{N,K}(u) + \lambda\cdot u$ over $u \in \mathcal{U}^{N,K}$. It is nonempty and compact by continuity and compactness. Let $D(\lambda) := \operatorname{diam} S_\lambda\Lambda^{-1}(\lambda)$, computed using the ℓ^1 norm. Every vector $S_\lambda u$ is nonnegative and has coordinates summing to at most $\overline{V}_N$, so $D(\lambda) \le 2\overline{V}_N$. Hence $M := \sup_{\lambda\in\mathbb{R}_{++}^N} D(\lambda) < \infty$. Fix $\lambda \in \mathbb{R}_{++}^N$. We show that any two promises in $\Lambda^{-1}(\lambda)$ have the same totals over the equal-costate groups.

Take $u, \widehat{u} \in \Lambda^{-1}(\lambda)$, and choose optimal policies at these states. Since $\lambda \in \mathbb{R}_{++}^N$, complementary slackness for promise keeping makes every worker's constraint bind under both policies.

For each realized v, write $m(v), z(v)$ and $\widehat{m}(v), \widehat{z}(v)$ for the two selected matchings and continuation promises. Lagrangian optimality makes both choices λ-optimal outside a common null set. Apply Lemma 7 after each such realization to rearrange each pair assortatively while preserving $S_\lambda G$. Group-level uniqueness gives both pairs the same task multiset within each equal-costate group. Relabel the workers together with their continuation coordinates by a fixed rule so that both pairs use the same canonical matching. These relabelings are measurable and preserve $S_\lambda G$. The pairs then have the same vector ℓ of required continuation promises and the same group-level current-utility vector. Summing the binding promise-keeping constraints within each group, the common current-utility term cancels, giving

$$S_\lambda u - S_\lambda \widehat{u} = \delta_A \mathbb{E}[S_\lambda z(v) - S_\lambda \widehat{z}(v)].$$

It remains to bound the difference between the continuation promises. Because the matching set is finite, the union of the exceptional sets in Lemma 8 is a null set. The lemma therefore applies simultaneously to every realized matching outside one null set.

Fix a realization outside this null set. Let z and $\widehat{z}$ solve the common continuation problem, and set $\bar{z} := (z + \widehat{z})/2$. The objective is concave and the feasible set is convex, so $\bar{z}$ is also a solution. Applying Lemma 8 at $\bar{z}$ gives $\alpha \in \mathbb{R}_+^N$ such that $\alpha_i(\bar{z}_i - \ell_i) = 0$ and $\mu := q\lambda + \alpha/\delta_P$ belongs to $\Lambda^{N,K}(\bar{z})$. Because $q > 0$ and $\lambda \in \mathbb{R}_{++}^N$, we have $\mu \in \mathbb{R}_{++}^N$. If $\alpha_i > 0$, then $\bar{z}_i = \ell_i$, and since $z_i, \widehat{z}_i \geq \ell_i$, both endpoints also equal ℓ_i in coordinate i.

The same costate μ supports both endpoints z and $\widehat{z}$. The continuation objective has the same value at z, $\bar{z}$, and $\widehat{z}$, while complementarity gives $\alpha \cdot (z - \bar{z}) = \alpha \cdot (\widehat{z} - \bar{z}) = 0$. Adding the multiplier term therefore leaves the three objective values equal. Since $\bar{z}$ maximizes this Lagrangian over $\mathcal{U}^{N,K}$, so do both endpoints. Hence $z, \widehat{z} \in \Lambda^{-1}(\mu)$.

Put $d := z - \widehat{z}$. Complementarity at $\bar{z}$ implies $d_i = 0$ whenever $\alpha_i > 0$. On every coordinate in the support of d, therefore, $\mu_i = q\lambda_i$. The partitions into equal μ-coordinates and equal λ-coordinates consequently coincide on the support of d; coordinates outside that support contribute zero. Thus $\|S_\lambda d\|_1 = \|S_\mu d\|_1 \leq D(\mu) \leq M$.

The preceding bound holds at almost every realization v. Taking norms in the promise-keeping identity and applying the triangle inequality gives $\|S_\lambda u - S_\lambda \widehat{u}\|_1 \leq \delta_A \mathbb{E}[\|S_\lambda(z(v) - \widehat{z}(v))\|_1] \leq \delta_A M$. Taking the supremum over $u, \widehat{u} \in \Lambda^{-1}(\lambda)$ gives $D(\lambda) \leq \delta_A M$. Taking the supremum over $\lambda \in \mathbb{R}_{++}^N$ gives $M \leq \delta_A M$. Since $\delta_A < 1$, $M = 0$. Thus all promises supported by λ have the same total in every equal-costate group. If λ has no ties, each group contains one worker and the entire promise vector is unique.

When $N = K = 1$, $\Lambda^{-1}(\lambda)$ contains one promise for every $\lambda > 0$. Any nontrivial affine segment in $(\bar{u}, V^+]$ has a strictly negative slope, whose negative is a common positive costate at multiple points of the segment. This contradicts uniqueness, so Φ is strictly concave on $(\bar{u}, V^+]$. □

Together, the proof of part (i) and Lemmas 5 and 9 prove Proposition 6. The specialization $N = K = 1$ gives Proposition 1.

A.3 Optimal Single-Agent Contract

This subsection proves Propositions 2 to 4 and Theorem 1. Use U and q from Section 3, and recall that the plateau $[0, \bar{u}]$ is the interval of promises on which Φ is constant. A supporting costate $\lambda \geq 0$ measures the marginal cost, in the Lagrangian, of delivering promised utility. Set $A_\lambda(z) := \delta_P \Phi(z) + \delta_A \lambda z$ and $B_\lambda(a) := \max_{z \in [a, V^+]} A_\lambda(z)$, with $B_\lambda(a) = -\infty$ if $a > V^+$. After optimizing the continuation promise, the difference between the Lagrangian value of assigning and rejecting a task of value v is

$$\Delta(\lambda; v) := (1 - \delta_P) + (1 - \delta_A)\lambda v + B_\lambda(\ell(v)) - B_\lambda(0).$$

Assignment is optimal when this gain is positive; either action is optimal at equality.

For each fixed λ, this gain is nondecreasing in v: its current term increases, while $\ell(v)$ decreases and therefore enlarges the feasible continuation set. If $\lambda > 0$, the gain is strictly increasing wherever the required continuation $\ell(v)$ is feasible. If $\lambda = 0$, it is positive when $\ell(v) \leq \bar{u}$, equals $-\infty$ when $\ell(v) > V^+$, and is strictly increasing between these regions because Φ is strictly decreasing above the plateau. Thus the assignment rule is a cutoff. The set of task values at which assignment and rejection are both optimal has probability zero in every case used below.

Proof of Proposition 2. Consider a static deterministic assignment rule $X : \mathcal{V} \to \{0, 1\}$, let $A := \{v : X(v) = 1\}$, and write $U_X := \mathbb{E}[vX(v)]$. Constraint (PC$_{ik}$) reduces to $X(v) = 1 \Rightarrow \ell(v) \leq U_X$.

Because F is continuous and strictly increasing on its support, there is a cutoff c_A satisfying $\Pr(V \geq c_A) = \Pr(V \in A)$. Among all measurable sets with this probability, the upper tail $[c_A, \overline{v}]$ maximizes expected task value, so $U(c_A) \geq U_X$. Constraint (PC$_{ik}$) also gives $A \subseteq \{v : \ell(v) \leq U_X\}$, and hence $\Pr(\ell(V) \leq U_X) \geq \Pr(V \in A) = \Pr(V \geq c_A)$. Since ℓ is nonincreasing and F is strictly increasing, $\ell(c_A) \leq U_X$. Therefore $\ell(c_A) \leq U(c_A)$, so the cutoff rule is feasible and has the same assignment probability as X.

An optimal cutoff cannot be positive. Lowering a positive cutoff to zero adds only positive-valued tasks, preserves (PC$_{ik}$), and strictly raises the assignment probability. It therefore remains to consider $c \in [\underline{v}, 0]$. For $c < 0$, feasibility is $U(c) \geq ((1 - \delta_A)/\delta_A)(-c)$, equivalently $H(c) \geq 0$. The function $H(c) = U(c) + ((1 - \delta_A)/\delta_A)c$ is strictly increasing on $[\underline{v}, 0]$, because $H'(c) = -cf(c) + (1 - \delta_A)/\delta_A > 0$, and $H(0) = V^+ > 0$. Thus the feasible cutoff set is all of $[\underline{v}, 0]$ when $H(\underline{v}) \geq 0$, and otherwise it is $[c^S, 0]$, where c^S is the unique root of $H(c) = 0$. Since the principal's assignment probability decreases with the cutoff, she chooses the lowest feasible cutoff. This gives the stated rule, stationary promise, and payoff. □

Proof of Proposition 3. Fix F. The static feasibility condition for a cutoff $c \le 0$ is $U(c) \ge ((1-\delta_A)/\delta_A)(-c)$. Increasing δ_A lowers the right side, so every cutoff that was feasible remains feasible. The firm chooses the lowest feasible cutoff; hence c^S weakly decreases and $\Pi^S = 1 - F(c^S)$ weakly increases. Lowering a nonpositive cutoff adds only weakly negative task values, so $U^S = U(c^S)$ weakly decreases. The principal's discount factor does not enter static feasibility, $U(c)$, or the stationary assignment probability.

Now let F_H first-order stochastically dominate F_L, let R be uniform on $(0, 1)$, and write $V_j = F_j^{-1}(R)$ for $j \in \{L, H\}$. Then $V_H \ge V_L$ almost surely. Transport the F_L-optimal rule by rank: under F_H, assign at rank R exactly when $X_L^S(V_L) = 1$. Since $R = F_H(V_H)$ almost surely, this is a measurable function of V_H and preserves the assignment probability. Its stationary utility $\widetilde{U}_H$ satisfies $\widetilde{U}_H \ge U_L^S$. At every assigned rank, $\ell(V_H) \le \ell(V_L) \le U_L^S \le \widetilde{U}_H$, so (PC$_{ik}$) holds. The transported rule is therefore feasible under F_H, and the optimal static assignment probability under F_H is weakly higher. □

Proof of Proposition 4. By Proposition 2, full static assignment is feasible if and only if $H(\underline{v}) \ge 0$, equivalently $\mathbb{E}[v] \ge ((1 - \delta_A)/\delta_A)(-\underline{v})$. In that case $c^S = \underline{v}$, and the principal obtains the maximal payoff one in every period.

Now suppose $c^S > \underline{v}$. Then $H(c^S) = 0$, so $U^S = U(c^S) = \ell(c^S)$. Because $c^S < 0$ and the density is positive, $U^S < V^+$. For all sufficiently small $\varepsilon > 0$, therefore, $c_\varepsilon := c^S - \varepsilon$ belongs to the support and $\ell(c_\varepsilon) < V^+$. From promise zero, assign exactly the tasks with $v \ge c_\varepsilon$, use continuation $u'(v) = \ell(v)$ on assigned negative values, and use continuation zero otherwise. These continuations are feasible, and (PC$_{ik}$) holds by construction. Delivered utility is nonnegative after every task realization: an assigned negative task delivers zero, an assigned nonnegative task delivers weakly positive utility, and a rejected task delivers zero. Its expectation is therefore nonnegative, so promise keeping at zero holds.

The static cutoff c^S can be run from every promise $u \le U^S$, so $\Phi(u) \ge \Pi^S$ on $[0, U^S]$. Outside the newly assigned interval $[c_\varepsilon, c^S)$, the policy just constructed uses a continuation no greater than U^S. On that interval, $\ell(v)$ converges uniformly to U^S as $\varepsilon \downarrow 0$. Continuity of Φ therefore implies that, for every $\eta > 0$ and all sufficiently small ε, $\Phi(\ell(v)) \ge \Pi^S - \eta$ throughout the new interval. Writing $m_\varepsilon := \Pr(c_\varepsilon \le v < c^S) > 0$, the current-payoff term is $(1 - \delta_P)(\Pi^S + m_\varepsilon)$, and the continuation term is at least $\delta_P(\Pi^S - \eta m_\varepsilon)$. The policy therefore has payoff at least $\Pi^S + (1 - \delta_P - \delta_P\eta)m_\varepsilon$. Choosing $\eta < (1 - \delta_P)/\delta_P$ proves $\Phi(0) > \Pi^S$.

It remains to prove finite-time exit from the plateau $\mathcal{P} = [0, \bar{u}]$, the interval on which the principal's value is constant. At every $u \in \mathcal{P}$, zero is a supporting costate, so pointwise Lagrangian optimality implies that, for almost every task value, an optimal decision maximizes the Bellman expression in which delivered utility has zero marginal cost. The gain $\Delta(0; v)$ is independent of u, and the preceding observation rules out ties except on a probability-zero set. All plateau states consequently use the same assignment rule almost everywhere.

The same pointwise optimality describes the continuation choice. The participation floor is $\ell(v)$ for an assigned negative task and zero otherwise. If this floor is at most $\bar{u}$, an optimal continuation remains on the plateau; choosing a continuation above $\bar{u}$ would strictly lower the continuation value after that realization. Hence an optimal decision can leave the plateau only after assigning a task with $\ell(v) > \bar{u}$, in which case (PC_{ik}) forces the next promise above the plateau.

Suppose, toward a contradiction, that an optimal decision at $\bar{u}$ kept its continuations on the plateau almost surely. Every continuation would then have value $\Phi(0)$, so the Bellman equation would give $\mathbb{E}[X(v)] = \Phi(0)$. Let $W := \mathbb{E}[vX(v)]$. Promise keeping and $u'(v) \leq \bar{u}$ imply $W \geq \bar{u}$, while $W \leq \mathbb{E}[[v]_+] = V^+$. Constraint ($\mathrm{PC}_{ik}$) gives $X(v) = 1 \Rightarrow \ell(v) \leq u'(v) \leq \bar{u} \leq W$. Thus the same assignment rule with constant continuation promise W is a feasible static contract: it delivers promise W exactly and satisfies (PC_{ik}). Its payoff is $\Phi(0)$, contradicting $\Phi(0) > \Pi^S$.

Thus every optimal decision at $\bar{u}$ exits with positive probability. At $\bar{u}$, exit occurs exactly after an assigned task whose participation floor exceeds $\bar{u}$: a floor at most $\bar{u}$ permits a continuation on the plateau, and choosing a continuation above the plateau would strictly reduce continuation value. Since assignment and rejection tie only on a probability-zero set, positive-probability exit implies that the set

$$E := \{v : \ell(v) \in (\bar{u}, V^+],\ \Delta(0; v) > 0\}$$

has positive probability. The gain $\Delta(0; v)$, and hence the assignment rule, is common across plateau states. Every optimal plateau decision therefore assigns almost every value in E, and (PC_{ik}) forces the next promise above $\bar{u}$. Set $p_E := \Pr(V \in E) > 0$. Conditional on remaining on the plateau, each period has exit probability at least p_E. Writing $\tau_{\mathcal{P}}$ for the first date outside $\mathcal{P}$, independence of task draws gives $\Pr(\tau_{\mathcal{P}} > n) \leq (1 - p_E)^n$, so exit occurs in finite time almost surely. □

Proof of Theorem 1. At the largest feasible promise $u = V^+$, the equality conditions in Lemma 4 force $vX(v) = [v]_+$ and $u'(v) = V^+$ almost surely. Thus this state assigns exactly the positive tasks and keeps continuation V^+.

Fix henceforth $u < V^+$, an optimal decision at that state, and a costate $\lambda \in \Lambda(u)$. Pointwise Lagrangian optimality implies that, after almost every realized v, the decision maximizes the Bellman expression augmented by λ times delivered utility. Assignment is therefore governed by the sign of $\Delta(\lambda; v)$, apart from a probability-zero set of ties.

At a fixed state $u < V^+$, every optimal decision is λ-optimal for every $\lambda \in \Lambda(u)$. For $\lambda > 0$, strict monotonicity of $\Delta(\lambda; \cdot)$ and the null tie set therefore pin down the same cutoff for every optimal decision and every positive supporting costate at that state. On the plateau, the argument with the zero costate gives the common cutoff γ_0. Since positive task values have strictly positive gain, every cutoff is nonpositive.

We next order the cutoffs above the plateau. If $\bar{u} < u_1 < u_2 < V^+$, $\lambda^{(j)} \in \Lambda(u_j)$, strict concavity gives $0 < \lambda^{(1)} < \lambda^{(2)}$. For a negative value v whose required continuation is feasible, set $a := \ell(v) > 0$. For $\lambda > 0$, let $r_\lambda \geq \bar{u}$ be the unique unconstrained maximizer of A_λ, and set $r_0 := \bar{u}$. A maximizer subject to $z \geq a$ is $\max\{a, r_\lambda\}$, and it is unique when $\lambda > 0$. Danskin's theorem gives

$$\frac{\partial \Delta(\lambda; v)}{\partial \lambda} = (1 - \delta_A)v + \delta_A (a - r_\lambda)_+,$$

with the same expression as the right derivative at zero. Since $\delta_A a = -(1 - \delta_A)v$, this derivative equals $(1 - \delta_A)v < 0$ when $a \leq r_\lambda$, and $-\delta_A r_\lambda < 0$ otherwise. Hence the gain from every feasible negative task strictly decreases with the costate. Since supporting costates strictly increase with the promise, the assigned sets shrink weakly as the promise rises, so the cutoff is nondecreasing, including at the plateau boundary. It may remain locally constant when the lowest dynamically feasible task pins it down, and $\gamma(V^+) = 0$.

We finally characterize continuation promises. After task value v is observed, let $a(v) := \ell(v)X(v; u)$ be the participation floor. If $u \in \mathcal{P}$, use the supporting costate zero. The maximizers of A_0 without the lower bound $z \geq a(v)$ are exactly the plateau states. If $a(v) \leq \bar{u}$, any point of $\mathcal{P} \cap [a(v), V^+]$ maximizes the continuation problem. If $a(v) > \bar{u}$, strict concavity and strict decline above the plateau make $a(v)$ the unique maximizer. A measurable choice of these value-contingent continuations must also deliver the promised utility in expectation before the task is drawn.

Now let $u \in (\bar{u}, V^+)$. Every positive task is assigned: giving a rejected positive task to the worker while retaining its continuation raises the principal's current payoff, relaxes promise keeping, and satisfies (PC_{ik}) without requiring a positive continuation promise. Fix $\lambda \in \Lambda(u)$. After almost every positive task value, (PC_{ik}) imposes no binding lower bound, so the continuation problem maximizes A_λ. Its first-order condition is $q\lambda \in \Lambda(z)$. Since $q\lambda > 0$, Proposition 1 implies that exactly one promise is compatible with this costate; call it t_λ.

The same optimal decision is supported by every costate in $\Lambda(u)$. On the positive-value set, which has positive probability, its continuation must therefore equal t_λ for every supporting costate. Comparing any two such costates outside their probability-zero exceptional sets shows that these promises coincide. Denote the common continuation target by $t(u)$. Equivalently, $q\Lambda(u) \subseteq \Lambda(t(u))$. Since every element of $q\Lambda(u)$ is strictly positive, $t(u) \geq \bar{u}$.

Because $q\lambda$ is finite while Lemma 4 rules out a finite supporting costate at V^+, the target satisfies $t(u) < V^+$.

This inclusion also determines whether the continuation target lies above or below the current promise. If $q > 1$, let $\lambda_{\max} := \max \Lambda(u)$. Then $q\lambda_{\max} \in \Lambda(t(u))$ and $q\lambda_{\max} > \lambda_{\max}$. Monotonicity of the supporting slopes of the convex function $-\Phi$ gives $t(u) \geq u$; equality is impossible because no costate at u exceeds $\lambda_{\max}$. Hence $t(u) > u$. If $q < 1$, let $\lambda_{\min} := \min \Lambda(u)$. Then $q\lambda_{\min} \in \Lambda(t(u))$ and $q\lambda_{\min} < \lambda_{\min}$, so the same monotonicity gives $t(u) \leq u$; equality is impossible because no costate at u lies below $\lambda_{\min}$. Hence $t(u) < u$. If $q = 1$, the same positive costate supports both states, and Lemma 9 says that it is compatible with only one promise, so $t(u) = u$. The target function is also nondecreasing. For $\bar{u} < u_1 < u_2 < V^+$, supporting costates satisfy $\lambda^{(1)} < \lambda^{(2)}$, so $q\lambda^{(1)} < q\lambda^{(2)}$; monotonicity of supporting slopes then gives $t(u_1) \leq t(u_2)$. Hence t is Borel measurable.

For each supporting costate $\lambda > 0$, A_λ is strictly increasing on the plateau and strictly concave above it, so its unconstrained maximizer is uniquely $t(u)$. Its unique maximizer subject to the lower bound $a(v)$ is therefore $\max\{a(v), t(u)\}$. Rejected values have $a(v) = 0$, while assigned values have $a(v) = \ell(v)$. This gives exactly the continuation rule in the theorem. Since assignment and rejection tie only on a probability-zero set, every optimal decision at a state above the plateau uses the same assignment and continuation rules almost everywhere. □

A.4 Optimal Multi-Agent, Multi-Task Contract

For the static optimality result, write $\mathcal{P} := \mathcal{P}^{N,K}$, $\Phi := \Phi^{N,K}$, and let $V^{\mathcal{P}}$ be the value of the Bellman problem whose continuation promises must remain in the plateau $\mathcal{P}$, where the principal's value is maximal.

Lemma 10 (Largest common promise on the plateau)**.** *Let $u^+ := \mathbb{E}[[V]_+]$, and define*

$$\rho_K := \max\{h \geq 0 : h\mathbf{1}_S \in \mathcal{P} \text{ for every } S \subseteq \{1, \dots, N\} \text{ with } |S| = K\}.$$

Then $\rho_K < u^+$. If $\mathbb{E}[V] < \ell(\underline{v})$, then $\rho_K < \ell(\underline{v})$.

Proof. The set in the definition contains zero and is closed. Every admissible h satisfies $Kh \leq \overline{V}_K = Ku^+$, so its maximum exists and $\rho_K \leq u^+$. For $r \leq K$, the sum of the r largest nonnegative task values is pointwise at least r/K times the sum of all K nonnegative task values. Hence $ru^+ \leq \overline{V}_r$. For any K-worker set S and any worker set T, put $s := |S \cap T|$. The total promise to T under $u^+\mathbf{1}_S$ is $su^+ \leq \overline{V}_s \leq \overline{V}_{|T|}$, so $u^+\mathbf{1}_S \in \mathcal{U}^{N,K}$. This vector has total promise $Ku^{+} = \overline{V}_N$. If it belonged to $\mathcal{P}$, the zero vector would support $-\Phi$ there, contrary to Lemma 4. Therefore $\rho_K < u^{+}$.

Now suppose $\mathbb{E}[V] < \ell(\underline{v})$ and $\rho_K \geq \ell(\underline{v})$. Fix a set S of K workers, assign the tasks bijectively to them, and give the recipient of task k continuation promise $\ell(V_k)$. Since $\ell(V_k) \leq \ell(\underline{v}) \leq \rho_K$, the continuation vector belongs to $\mathcal{P}$ by downward closure. The decision satisfies (PC$_{ik}$) and delivers nonnegative utility to every worker, so it is feasible from promise zero. It gives $\Phi(0) \geq (1-\delta_P)K + \delta_P\Phi(0)$; the upper bound $\Phi(0) \leq K$ therefore implies $\Phi(0) = K$.

Let $M := \max\{\sum_i u_i : \Phi(u) = K\}$, which exists by compactness. Since $\ell(\underline{v})\mathbf{1}_S \in \mathcal{P}$, we have $M \geq K\ell(\underline{v})$. At a state attaining M, Bellman optimality and the upper bound K force an optimal policy to assign all K tasks and choose a continuation z satisfying $\Phi(z) = K$ almost surely. Hence $\sum_i z_i \leq M$ almost surely. Summing promise keeping gives $M \leq (1-\delta_A)K\mathbb{E}[V] + \delta_A M$, and therefore $M \leq K\mathbb{E}[V] < K\ell(\underline{v})$, a contradiction. □

Proof of Proposition 8. At most K tasks can be assigned in any period, so K is an upper bound on the principal's payoff in any dynamic contract. Suppose first that $\mathbb{E}[V] \geq \ell(\underline{v})$. Assign task k to agent k in every period, set the first K agents' continuation promises equal to $\mathbb{E}[V]$, and set all other promises to zero. Promise keeping holds with equality. The promise vector lies in $\mathcal{U}^{N,K}$: for $r \leq K$, the sum of the best r nonnegative task values dominates the sum of any fixed r task values, so $r\mathbb{E}[V] \leq \overline{V}_r$; a subset with more than K agents contains at most all K nonzero coordinates, so its total promise is at most $K\mathbb{E}[V] \leq \overline{V}_K = \overline{V}_r$. Constraint (PC$_{ik}$) holds because every assigned value $x \geq \underline{v}$ satisfies $(1-\delta_A)x + \delta_A\mathbb{E}[V] \geq (1-\delta_A)\underline{v} + \delta_A\ell(\underline{v}) = 0$. This contract assigns all K tasks every period and attains the dynamic upper bound.

Conversely, suppose a feasible static contract with fixed promise vector u assigns all K tasks after almost every realization. Summing promise keeping across agents gives $\sum_i u_i \leq K\mathbb{E}[V]$. Fix $\eta > 0$. On the positive-probability event that all task values lie in $[\underline{v}, \underline{v}+\eta]$, only agents with $u_i \geq \ell(\underline{v}+\eta)$ can receive positive assignment probability. Each such agent can receive total assignment probability at most one, whereas full assignment has total probability K. There must therefore be at least K such agents, so $\sum_i u_i \geq K\ell(\underline{v}+\eta)$. Letting $\eta \downarrow 0$ gives $\mathbb{E}[V] \geq \ell(\underline{v})$. Thus full static assignment is feasible exactly under the stated condition.

Now assume $\mathbb{E}[V] < \ell(\underline{v})$. I first show that $\Phi(0) > V^{\mathcal{P}}(0)$. By Lemma 10, $\rho_K < \min\{u^{+}, \ell(\underline{v})\}$; coordinatewise downward closure also gives $h\mathbf{1}_S \in \mathcal{P}$ whenever $h \leq \rho_K$ and $|S| = K$.

Define the uniform-continuity bound $\omega(\eta) := \sup\{|\Phi(x) - \Phi(y)| : x, y \in \mathcal{U}^{N,K},\ \|x-y\|_\infty \leq \eta\}$. Choose $\eta > 0$ with $\rho_K + \eta < \min\{u^{+}, \ell(\underline{v})\}$ and $\delta_P\omega(\eta) < 1 - \delta_P$. Let $E_\eta := \{\ell(V_k) \in (\rho_K, \rho_K + \eta)$ for every $k = 1, \ldots, K\}$. This event has positive probability. On E_η, every task is negative and requires a continuation promise above ρ_K. A policy whose continuation promises must remain in $\mathcal{P}$ can assign at most $K-1$ tasks there. If it assigned all K tasks to a set S and chose $z \in \mathcal{P}$, then $h := \min_{i \in S} z_i > \rho_K$. Downward closure gives $h\mathbf{1}_S \in \mathcal{P}$, and symmetry gives $h\mathbf{1}_T \in \mathcal{P}$ for every K-worker set T, contradicting the definition of ρ_K.

Take a plateau-restricted policy from 0 whose payoff is at least $V^{\mathcal{P}}(0) - \xi$. Modify its date-zero decision only on E_η: assign all K tasks to distinct agents and set each assigned agent's continuation promise equal to her participation floor, with every other continuation promise equal to zero. Continue optimally after this decision. The modified continuation vector is feasible because it is dominated by the feasible vector $u^+\mathbf{1}_S$ for the assigned set S. It is also within distance η of $\rho_K\mathbf{1}_S \in \mathcal{P}$, measured by the largest coordinate difference, so its continuation value is at least $\Phi(0) - \omega(\eta)$. The modification preserves promise keeping at 0, since each assigned negative task with continuation promise $\ell(v)$ gives total utility $(1-\delta_A)v + \delta_A\ell(v) = 0$, and unassigned agents receive zero. Conditional on E_η, the original policy assigns at most $K-1$ tasks and has continuation value at most $\Phi(0)$, while the modified policy assigns K tasks and has continuation value at least $\Phi(0) - \omega(\eta)$. Its conditional payoff gain is therefore at least $g := (1-\delta_P) - \delta_P\omega(\eta) > 0$. The modified payoff is at least $V^{\mathcal{P}}(0) - \xi + \Pr(E_\eta)g$. Choosing $\xi < \Pr(E_\eta)g$ gives $\Phi(0) > V^{\mathcal{P}}(0)$.

Now let Π^S be the optimal static payoff. Every static contract is feasible from promise zero, so $\Pi^S \leq \Phi(0)$. Suppose $\Pi^S = \Phi(0)$, and let u^S be the fixed promise vector of an optimal static contract. That contract is feasible from u^S, giving $\Phi(u^S) \geq \Pi^S = \Phi(0)$. Monotonicity gives the reverse inequality, so $u^S \in \mathcal{P}$. Starting from promise zero, the same static contract with constant continuation u^S keeps every continuation in $\mathcal{P}$, implying $\Pi^S \leq V^{\mathcal{P}}(0) < \Phi(0)$, a contradiction. Therefore $\Phi^{N,K}(0) > \Pi^S$.

The final claim now follows by cases. When static full assignment is feasible, it attains the upper bound K and is dynamically optimal from promise zero. When it is infeasible, the strict comparison just proved rules out every static contract from being dynamically optimal from zero. Thus a static contract is dynamically optimal from promise zero exactly when it assigns all K tasks after almost every realization. □

Proof of Proposition 9. We prove the three comparisons in turn.

Worker patience. Fix b and F, write a for the workers' common discount factor, and set $\ell_a(v) := [-(1-a)v/a]_+$. The promise polytope $\mathcal{U}^{N,K}$ does not depend on a and is coordinatewise downward closed. Let $a_2 > a_1$ and $\rho := (1-a_2)/(1-a_1)$. Take any contract feasible under a_1 from u, retain every matching, and multiply every current and continuation promise by ρ. If Y_i is worker i's current task value and z_i is her old continuation promise, then

$$\mathbb{E}[(1-a_2)Y_i + a_2\rho z_i] = \rho\mathbb{E}[(1-a_1)Y_i + a_1 z_i] + \rho(a_2 - a_1)\mathbb{E}[z_i] \geq \rho u_i.$$

If worker i receives a negative task of value v, constraint (PC_{ik}) under a_1 gives

$$\rho z_i \geq \frac{1-a_2}{a_1}(-v) \geq \frac{1-a_2}{a_2}(-v),$$

so (PC_{ik}) also holds under a_2. Downward closure keeps every scaled promise feasible. The assignment process and the principal's payoff are unchanged, so $\Phi^{N,K}_{a_2}(\rho u) \geq \Phi^{N,K}_{a_1}(u)$. Setting $u = 0$ proves the weak payoff comparison.

For strict monotonicity before full assignment, use unnormalized promises $w = u/(1-a)$, let $\mathcal{W}_a := (1-a)^{-1}\mathcal{U}^{N,K}$, and write $\Psi_a(w) := \Phi^{N,K}_a((1-a)w)$. In these units the constraints are $\mathbb{E}[Y_i + aw'_i] \geq w_i$ and $v_k + aw'_i \geq 0$ when worker i receives task k. Since $a_2 > a_1$, $\mathcal{W}_{a_1} \subseteq \mathcal{W}_{a_2}$, and the preceding comparison gives $\Psi_{a_2} \geq \Psi_{a_1}$ on $\mathcal{W}_{a_1}$.

Suppose $\Psi_{a_1}(0) < K$. Static full assignment is then infeasible, so the plateau-restricted gap established in the proof of Proposition 8 applies. Let $\mathcal{Q}$ be the plateau of Ψ_{a_1} and follow an a_1-optimal law from zero. Define its first exit time by $\tau := \inf\{t \geq 0 : w_{t+1} \notin \mathcal{Q}\}$. The plateau-restricted gap implies $\Pr(\tau < \infty) > 0$. Under a_2, follow the old law while its continuation remains in $\mathcal{Q}$. At the first history and task realization for which the old continuation $w'(v)$ lies outside $\mathcal{Q}$, set $c := a_1/a_2$ and replace that continuation by $cw'(v)$; retain $w'(v)$ on realizations that remain in $\mathcal{Q}$. Downward closure gives $cw'(v) \in \mathcal{W}_{a_1} \subseteq \mathcal{W}_{a_2}$. Promise keeping and ($\mathrm{PC}_{ik}$) remain satisfied because $a_2cw' = a_1w'$ on the modified realizations and $a_2w' \geq a_1w'$ on the rest. On histories for which $\tau = \infty$, follow the old law forever.

For every modified continuation, concavity and $w' \notin \mathcal{Q}$ give

$$\Psi_{a_1}(cw') \geq (1-c)\Psi_{a_1}(0) + c\Psi_{a_1}(w') > \Psi_{a_1}(w').$$

Moreover, $\Psi_{a_2}(cw') \geq \Psi_{a_1}(cw')$. Continuing optimally under a_2 therefore preserves the old payoff before this transition and raises the continuation payoff on an event of positive probability. The expected discounted gain is strictly positive, proving $P^{N,K}(a_2, b, F) > P^{N,K}(a_1, b, F)$.

For the limit as $a \downarrow 0$, every individual continuation promise is at most $\overline{V}_1$. Constraint (PC_{ik}) therefore requires any assigned negative value to satisfy $v \geq -a\overline{V}_1/(1-a)$. At every history, expected assignments are bounded above by $K \Pr(V \geq -a\overline{V}_1/(1-a))$. Assigning all positive tasks supplies the lower bound $K \Pr(V > 0)$, so the density assumption and the squeeze theorem prove the first limit.

For the limit as $a \uparrow 1$, fix an integer T, put $u^+ := \mathbb{E}[V_+]$, and choose a fixed set S of K workers. For $h \in [0, u^+]$, the vector $h\mathbf{1}_S$ belongs to $\mathcal{U}^{N,K}$. Indeed, if a subset of r workers contains s members of S, its total promise is at most $su^+ \leq \overline{V}_s \leq \overline{V}_r$; the first inequality follows because the sum of the s largest nonnegative task values is pointwise at least s/K times the sum of all K nonnegative values.

Set $h_0 := u^+$ and $h_n := (1-a)\mathbb{E}[V] + ah_{n-1} = (1-a^n)\mathbb{E}[V] + a^n u^+$ for $n = 1, \ldots, T$. When a is sufficiently close to one, every h_n is nonnegative and $(1-a)\underline{v} + ah_{n-1} \geq 0$. Fix a bijection between the K tasks and the workers in S. From promise zero, assign all K tasks using this matching and choose continuation $h_{T-1}\mathbf{1}_S$; this delivers $h_T\mathbf{1}_S \geq 0$ in expectation. At each subsequent state $h_n\mathbf{1}_S$, for $n = T-1, \ldots, 1$, use the same matching with continuation $h_{n-1}\mathbf{1}_S$; this delivers $h_n\mathbf{1}_S$ in expectation, and the preceding worst-task inequality ensures (PC_{ik}). After T full-assignment dates, assign only positive tasks and keep promise $u^+\mathbf{1}_S$. This rule is feasible from zero and gives payoff $K\{1 - b^T \Pr(V < 0)\}$. First choose T large and then take a close to one. Since K is an upper bound, the second limit follows.

Principal patience. Fix a and F, suppress them from the notation, and let $b_1 < b_2$. The feasible contracts do not depend on b. Follow a b_2-optimal Markov contract from zero, let M_t be its number of assignments at date t, and define

$$A_t := \mathbb{E}[M_t], \qquad B_t := \mathbb{E}[\Phi_{b_2}^{N,K}(u_t)].$$

Bellman optimality gives $B_t = (1-b_2)A_t + b_2 B_{t+1}$. Coordinatewise monotonicity gives $B_t \leq B_0 = \Phi_{b_2}^{N,K}(0)$. Evaluating the same contract at b_1 and rearranging the absolutely convergent sums yields

$$(1-b_1)\sum_{t\geq 0} b_1^t A_t - B_0 = \frac{(b_2-b_1)(1-b_1)}{1-b_2}\sum_{t\geq 1} b_1^{t-1}(B_0 - B_t) \geq 0.$$

Optimization at b_1 proves $P^{N,K}(a, b_1, F) \geq P^{N,K}(a, b_2, F)$.

If static full assignment is infeasible and every B_t equalled B_0, then u_t would belong to the b_2-plateau almost surely at every finite date. The countable intersection of these probability-one events would make this optimal contract feasible for the plateau-restricted problem, contradicting the strict restricted-value comparison in the proof of Proposition 8. Hence $B_t < B_0$ for some finite t, making the comparison strict. Under full assignment, both values equal K.

The unnormalized discounted task count is

$$\widehat{P}^{N,K}(b) := \frac{P^{N,K}(a,b,F)}{1-b} = \sup_{\pi} \mathbb{E}_{\pi}\left[\sum_{t\geq 0} b^t M_t\right].$$

Choose a b_1-optimal contract with the task-selection property in Theorem 2. It assigns every positive task at every date, so $\mathbb{E}[M_1] \geq K\Pr(V > 0) > 0$. Evaluating this contract at b_2 strictly raises its unnormalized sum, and optimization proves $\widehat{P}^{N,K}(b_2) > \widehat{P}^{N,K}(b_1)$. A static contract has the same expected assignment count in every period, so its normalized payoff is independent of b. This also proves the claim about the gain from dynamic contracting. Finally, if T_b is independent of $(M_t)_{t\geq 0}$ and satisfies $\Pr(T_b = t) = (1-b)b^t$, then

$$\mathbb{E}\left[\frac{M_{T_b}}{K}\right] = \frac{1-b}{K}\mathbb{E}\left[\sum_{t\geq 0} b^t M_t\right].$$

Applying this identity to a b-optimal contract gives the interpretation in the text.

Task values. Let F_H first-order stochastically dominate F_L. Couple every task by independent uniform ranks R_k and set $V_{j,k} := F_j^{-1}(R_k)$ for $j \in \{L, H\}$. Then $V_{H,k} \geq V_{L,k}$ almost surely. The ordered positive values, and hence every bound $\overline{V}_m$, are weakly higher under F_H. Therefore $\mathcal{U}_{F_L}^{N,K} \subseteq \mathcal{U}_{F_H}^{N,K}$.

Take any contract feasible under F_L. In the F_H economy, after observing $V_{H,k}$, set $R_k = F_H(V_{H,k})$ and $\widetilde{V}_{L,k} = F_L^{-1}(R_k)$, apply the old decision rule to $\widetilde{\mathbf{V}}_L$, and retain its continuation promise. The standing density assumption makes the ranks independent and uniform, and every old continuation remains feasible because $\mathcal{U}_{F_L}^{N,K} \subseteq \mathcal{U}_{F_H}^{N,K}$. Every assigned worker's current utility is weakly higher, every participation floor is weakly lower, and promise keeping is relaxed. The assignment and continuation-promise processes have the same law as before, so the principal obtains the same payoff. Taking suprema proves the value comparison.

When $N = K = 1$, the promise set and Bellman problem coincide with their scalar counterparts. The randomized static benchmark used in the multi-worker formulation also has the same value as the deterministic scalar benchmark. Indeed, Lemma 3 preserves the assignment probability and expected task utility; using the latter as the stationary promise gives a feasible deterministic scalar rule, whose optimum is characterized in Proposition 2. Thus the proposition also gives the comparative statics of the single-worker contract.

□

An equal-costate group contains workers whose costate coordinates are equal. Within such a group, the linear part of the Lagrangian depends only on the group's total delivered utility. Relabeling tasks within the group can redistribute utility across its workers without changing that term, although it can affect their individual promise-keeping constraints. The next lemma resolves this conflict with a public lottery: before the task vector is drawn, a random outcome observed by everyone selects a promise state and a decision that is optimal at that state.

We will use optimal deterministic one-period decisions with the following four properties at a state $x \in \mathcal{U}^{N,K}$. First, the decision delivers exactly the promised utilities: $\mathbb{E}[(1-\delta_A)Y + \delta_A z] = x$. Second, it is promise-assortative at x: after each idle worker is represented by a dummy task of value zero, a worker with a higher promise receives a weakly higher task value. Third, it assigns every positive task. Fourth, if it assigns a negative task, it also assigns every realized negative task with a higher value.

Lemma 11 (Assortative promise splitting)**.** *Suppose $u \in \mathcal{U}^{N,K}$ has a finite supporting costate. There are finitely many states $x^1, \dots, x^A$, each of which admits an optimal deterministic one-period decision with the four properties stated above, and probabilities $p_1, \dots, p_A$ such that*

$$\sum_a p_a x^a \geq u \qquad \textit{and} \qquad \sum_a p_a \Phi^{N,K}(x^a) = \Phi^{N,K}(u).$$

If u has a strictly positive supporting costate, the first relation can be chosen as equality.

Proof. Write $\Phi := \Phi^{N,K}$, $\mathcal{U} := \mathcal{U}^{N,K}$, and first fix u with a finite costate $\lambda \in \Lambda^{N,K}(u)$. Let

$$E_\lambda := \Lambda^{-1}(\lambda) = \arg\max_{x \in \mathcal{U}} \{\Phi(x) + \lambda \cdot x\}.$$

Thus E_λ is the nonempty compact convex set of all promise vectors supported by the same costate λ; it contains u, and $\Phi(x) + \lambda \cdot x$ is constant on it. Let $P := \{i : \lambda_i > 0\}$, $Z := \{i : \lambda_i = 0\}$, and define

$$D := \{r \in \mathbb{R}^N : r_i \geq 0 \text{ for every } i \in Z\}.$$

The restriction $r_i \geq 0$ on Z ensures that $\lambda + \varepsilon r$ remains nonnegative for small positive ε. Choose r in the interior of D so that $r \cdot x$ has a unique maximizer over E_λ and the coordinates of r differ within every equal-costate group. Denote that unique maximizer by $s(r)$. For sufficiently small $\varepsilon > 0$, the perturbed costate $\lambda + \varepsilon r$ is strictly positive and has distinct coordinates. Lemma 9 therefore gives a unique promise vector supported by it; call this vector u^ε. The defining optimality inequalities imply that $u^\varepsilon \to s(r)$ as $\varepsilon \downarrow 0$. Indeed, every limit point first maximizes $\Phi(x) + \lambda \cdot x$, and comparison of the terms multiplied by ε then makes it maximize $r \cdot x$ on E_λ.

Apply Lemma 6 to a Bellman-optimal policy at u^ε, using the supporting costate $\lambda + \varepsilon r$. This gives a measurable optimal policy with the stated task-selection properties and with assignments ordered by that costate. Because the costate is strictly positive, complementary slackness makes every promise-keeping constraint bind. If $u_i^\varepsilon < u_j^\varepsilon$, the Schur–Ostrowski inequality gives $(\lambda + \varepsilon r)_i \leq (\lambda + \varepsilon r)_j$; distinctness makes the inequality strict, so costate assortativity gives $Y_i \leq Y_j$. The policy is therefore promise-assortative. View these policies as probability laws on the compact space of task vectors, matchings, and continuation promises. Along a convergent subsequence, their limit is feasible and optimal at $s(r)$, delivers $s(r)$, and remains promise-assortative because $(x_i - x_j)(Y_i - Y_j) \geq 0$ defines a closed set. Every strict comparison in the original λ persists for small ε, so the limit also satisfies $\lambda_i < \lambda_j \Rightarrow Y_i \leq Y_j$. The task-selection properties pass to the limit because task ties and zero values have probability zero.

The limit may initially randomize among one-period choices after observing the task vector. For each matching used with positive conditional probability after v, replace its continuation lottery by $\mathbb{E}[z \mid v, m]$. Because $\mathcal{Z}(m, v)$ is convex, this conditional mean remains feasible for matching m. Concavity weakly raises the continuation value, and optimality makes this increase zero. Apply Lemma 3 to the resulting finite lottery over matchings, preserving the principal's payoff and each worker's delivered utility and restricting the deterministic rule to matchings used by the lottery. This preserves promise assortativity and task selection, so $s(r)$ admits an optimal deterministic decision with all four properties.

To verify that the existence of such a decision is preserved under limits, let $\mathcal{Q}(x)$ be the set of probability laws over one-period choices, allowing the randomization just described, that satisfy matching feasibility and (PC$_{ik}$), deliver x exactly, attain payoff $\Phi(x)$, and assign probability one to promise-assortative matchings with the stated task-selection properties. The compactness argument in Appendix A.1 and continuity of Φ imply that the graph $\{(x, Q) : Q \in \mathcal{Q}(x)\}$ is closed in the compact product of the promise set and the space of one-period laws. The support restrictions are closed: the ordering condition is $(x_i - x_j)(Y_i - Y_j) \geq 0$, every positive task is matched, and no assigned negative task has a lower value than a rejected negative task. The graph is therefore compact. The preceding deterministic-replacement argument shows that its projection onto the promise set is exactly the set of states that admit such a decision, so this set is compact.

Let $\mathcal{S}_\lambda$ be the closure of the states $s(r)$ obtained from perturbations satisfying the preceding uniqueness and coordinate conditions, and let $C_\lambda := \mathrm{co}(\mathcal{S}_\lambda)$ be their convex hull. The preceding limit argument shows that every state in $\mathcal{S}_\lambda$ continues to admit an optimal deterministic decision with the four stated properties. Compactness of the associated decisions also preserves the ordering across distinct equal-costate groups: $\lambda_i < \lambda_j \Rightarrow Y_i \leq Y_j$. If worker labels are permuted within an equal-costate group, the construction permutes the resulting state and decision in the same way. Thus $\mathcal{S}_\lambda$, and hence C_λ, is unchanged by such within-group permutations. Finally, $\mathcal{S}_\lambda \subseteq E_\lambda$, so $C_\lambda \subseteq E_\lambda$. For a compact set C, its support function in direction r is the largest value of $r \cdot x$ over $x \in C$. The support functions of C_λ and E_λ agree for every $r \in D$:

$$\max_{x \in C_\lambda} r \cdot x = \max_{x \in E_\lambda} r \cdot x, \qquad r \in D.$$

To see this, the support function of a compact convex set is convex and is differentiable on a dense subset of every open set; differentiability here is equivalent to having a unique maximizer. Thus directions with a unique maximizer are dense in the interior of D, and the requirement that the coordinates of r differ within equal-costate groups can also be imposed on a dense subset. Equality holds at every such direction by construction. Both support functions are continuous, so equality extends to all of D, including its boundary.

We next show that C_λ contains a vector y satisfying $y_i = u_i$ for $i \in P$ and $y_i \geq u_i$ for $i \in Z$. If it did not, the separating-hyperplane theorem would give a strict linear separation between the compact convex set C_λ and

$$A_u := \{y \in \mathbb{R}^N : y_i = u_i \text{ for } i \in P,\ y_i \geq u_i \text{ for } i \in Z\}.$$

Thus there would be a vector r such that $\max_{x \in C_\lambda} r \cdot x < \inf_{y \in A_u} r \cdot y$. The infimum on the right is finite only if $r_i \geq 0$ on Z; hence $r \in D$, and the infimum equals $r \cdot u$. Since $u \in E_\lambda$, this contradicts equality of the two support functions on D.

By Carathéodory's theorem, write this vector as $y = \sum_a p_a x^a$ using finitely many states from $\mathcal{S}_\lambda$. Then $y \geq u$ and $\lambda \cdot y = \lambda \cdot u$. If c_λ is the common value of $\Phi(x) + \lambda \cdot x$ on E_λ,

$$\sum_a p_a \Phi(x^a) = c_\lambda - \lambda \cdot y = c_\lambda - \lambda \cdot u = \Phi(u).$$

When $\lambda \gg 0$, the set Z is empty, so the construction gives $y = u$. □

Proof of Theorem 2. Fix a promise state u. We first construct a public lottery whose outcome a selects a promise state x^a and a one-period decision that is optimal at x^a. The resulting lottery is feasible and optimal at u, and its expected task values respect the order of u. Writing Y_i for worker i's task value and setting $Y_i = 0$ when she is idle, the construction gives

$$(u_i - u_j)\left(\sum_a p_a \mathbb{E}[Y_i \mid a] - \sum_a p_a \mathbb{E}[Y_j \mid a]\right) \geq 0 \qquad \text{for every } i, j.$$

We then eliminate the lottery while preserving these expected values and the two task-selection properties in the theorem. This produces the deterministic Markov decision asserted in the theorem.

First suppose u has a finite supporting costate λ. Retain the sets $\mathcal{S}_\lambda \subseteq E_\lambda$ and $C_\lambda = \mathrm{co}(\mathcal{S}_\lambda)$ from the proof of Lemma 11. That proof shows that C_λ is invariant under permutations within equal-costate groups and that every state in $\mathcal{S}_\lambda$ admits an optimal deterministic decision with the four stated properties and with $\lambda_i < \lambda_j \Rightarrow Y_i \leq Y_j$ after almost every task realization. It also produces a vector $y \in C_\lambda$ such that $y_i = u_i$ when $\lambda_i > 0$, $y_i \geq u_i$ when $\lambda_i = 0$, and $\lambda \cdot y = \lambda \cdot u$. Within the zero-costate group, permute the coordinates of y into the same order as those of u. The new vector remains in C_λ, and it still satisfies $y \geq u$ because the ordered coordinates of y dominate the corresponding ordered coordinates of u. In every positive-costate group, y already equals u.

The remaining construction has three steps. We sort the selected promise state after every lottery outcome, record the resulting mean w, and then use an order-preserving doubly stochastic matrix to restore the target mean y without reversing expected assignment quality.

Choose a finite representation $y = \sum_a p_a \widetilde{x}^a$ with $\widetilde{x}^a \in \mathcal{S}_\lambda$, and attach to each state one of these optimal deterministic decisions, retaining the strict ordering across groups with different costates. Fix an equal-costate group B, temporarily label its workers in increasing order of u_i, and relabel each state and decision within B so that the relabelled vector x_B^a is increasing. If several coordinates of x_B^a are equal, randomize uniformly over permutations of those workers. This leaves the promise vector unchanged and gives tied workers identical assignment distributions. Promise assortativity then makes $\mathbb{E}[Y_B \mid a]$ increasing, including at promise ties. This finite random relabelling can be included as part of the public lottery after the following averages are formed. Define

$$w_B := \sum_a p_a x_B^a, \qquad c_B := \sum_a p_a \mathbb{E}[Y_B \mid a]$$

after these outcome-specific relabellings. Both vectors are increasing. Let w and c denote the full vectors whose restrictions to each group are w_B and c_B.

The target y_B is majorized by w_B. Indeed, the sum of the largest r coordinates is a convex function, so its value at $y_B = \sum_a p_a \widetilde{x}_B^a$ cannot exceed the average sum of the largest r coordinates of the original promise states. Sorting does not change those sums, and their average is the sum of the largest r coordinates of w_B. The two vectors have the same total. The standard extreme-point description of increasing vectors majorized by w_B writes y_B as a convex combination of vectors obtained by averaging w_B over consecutive blocks (see Marshall, Olkin, & Arnold, 2011). The same convex combination of the corresponding block-averaging matrices gives a doubly stochastic matrix H_B such that $H_B w_B = y_B$. Each block-averaging matrix preserves increasing vectors, so $H_B c_B$ is increasing.

Do this independently in every equal-costate group and write H for the resulting block-diagonal doubly stochastic matrix. By the Birkhoff–von Neumann theorem, H is the mean of a finite lottery over permutations within those groups. Apply that same permutation lottery to worker labels after every original lottery outcome, permuting the promise state, matching, and continuation promise together. The new mean promise is $Hw = y$, while its expected task-value vector is Hc.

This vector has the required order. Within an equal-costate group, $H_B c_B$ is increasing in the order of u. Across groups with different costates, the Schur–Ostrowski criterion gives $u_i < u_j \Rightarrow \lambda_i \le \lambda_j$; the inequality is strict because the workers belong to different costate groups. The retained assignment ordering then gives $Y_i \le Y_j$ after every lottery outcome and task realization. Thus the displayed average-order inequality holds for every pair of workers.

Each attached decision delivers its promise state exactly, so the lottery delivers its mean $y \geq u$. Every relabelled state remains in E_λ, and $\lambda \cdot y = \lambda \cdot u$, so its expected Bellman payoff is $\Phi(u)$. The lottery is therefore feasible and optimal at u. Relabelling preserves promise assortativity after every outcome and the task-selection conclusions.

For a state without a finite supporting costate, choose a symmetric point $b = \beta\mathbf{1}$ in the relative interior of $\mathcal{U}$ and let $u^n := (1 - 1/n)u + (1/n)b$. This approximation preserves every coordinate ordering in u, because $u^n_i - u^n_j = (1 - 1/n)(u_i - u_j)$. Since u^n lies in the relative interior, the convex function $-\Phi$ has a finite subgradient there; coordinatewise monotonicity makes every such subgradient nonnegative. Apply the finite-costate construction at u^n and regard each resulting lottery as a probability measure on the compact graph $\{(x, Q) : Q \in \mathcal{Q}(x)\}$ defined in the preceding lemma's proof. A subsequence converges in distribution to a lottery whose mean promise dominates u, whose expected Bellman payoff is $\Phi(u)$, and whose expected task-value vector is ordered by u. These conclusions follow from continuity of Φ, preservation of membership in $\mathcal{Q}(x)$ under limits, and continuity of expected current task values. In detail, $Y_i(m, v)$ is bounded and continuous when the finite matching set is given the discrete topology, so $Q \mapsto \mathbb{E}_Q[Y_i]$ is continuous under convergence in distribution.

Apply Carathéodory's theorem to the image of this lottery under $(x, Q) \mapsto (x, \Phi(x), \mathbb{E}_Q[Y])$. Its mean can be represented using finitely many outcomes, preserving the mean promise, Bellman payoff, and expected task-value vector.

For each u, view the public lottery and its attached decisions together as one randomized one-period decision. It is feasible, attains $\Phi(u)$, and has the ordered expected task values just constructed. The correspondence of randomized decisions with these properties has nonempty compact values and a closed graph: feasibility and the ordering and task-selection restrictions are closed, while optimality follows from continuity of Φ. A measurable selection therefore exists. Apply Lemma 3 to this selection, preserving the expected coordinates of Y and restricting the deterministic choice to matchings used by the lottery. The measurable-selection argument in the proof of that lemma makes the resulting assignment and continuation rules jointly measurable in the promise state and task vector. Applying this Bellman-optimal rule recursively from the initial promise gives the asserted optimal deterministic Markov contract. □

The assignment result characterizes the current matching. A full recursive contract also requires continuation promises. The next proposition applies the standard first-order conditions for this constrained continuation choice.

Proposition 10 (Continuation costate equation)**.** *There is an F^K-null set E such that, for every state $u \in \mathcal{U}^{N,K}$,* *every $\lambda \in \Lambda^{N,K}(u)$, every $v \notin E$, and every (m, z) that is λ-optimal at v, there exist a multiplier $\mu \in \mathbb{R}_+^N$ on (PC$_{ik}$) and a next-period costate $\lambda^+ \in \Lambda^{N,K}(z)$ satisfying $\mu_i(z_i - \ell_i(m, v)) = 0$ for every i and $\lambda^+ = q\lambda + \mu/\delta_P$, where $q = \delta_A/\delta_P$. Thus μ_i can be positive only when (PC$_{ik}$) binds for worker i. The multiplier and next-period costate may each be nonunique.*

Proof. For each matching m, Lemma 8 holds outside an F^K-null set E_m. The matching set is finite, so $E := \bigcup_m E_m$ is an F^K-null set. Fix u, $\lambda \in \Lambda^{N,K}(u)$, $v \notin E$, and a λ-optimal (m, z). By definition, z attains $\Gamma_\lambda(\ell(m, v))$. The lemma therefore gives $\mu \in \mathbb{R}_+^N$ and $\lambda^+ \in \Lambda^{N,K}(z)$ satisfying the complementarity conditions and $\lambda^+ = q\lambda + \mu/\delta_P$. □

A.5 Long-Run Multi-Agent Dynamics

This subsection proves Theorem 3 and derives its implications for the long-run fraction of tasks assigned; its final part specializes the argument to prove Proposition 5. Throughout, write $\Phi := \Phi^{N,K}$, $\mathcal{U} := \mathcal{U}^{N,K}$, $\mathcal{P} := \mathcal{P}^{N,K}$, and $\Lambda(u) := \Lambda^{N,K}(u)$. Recall that $\mathcal{P}$ is the *plateau*: the promise vectors at which the principal obtains the same maximal value as at promise zero. Let $\bar{h} := \ell(\underline{v})$, let $\bar{f} := \sup_{v \in [\underline{v}, \overline{v}]} f(v) < \infty$, and let $\mathcal{F}_t$ be the history immediately before the date-t task draw, including any public randomization conducted before that draw. Thus V_t is independent of $\mathcal{F}_t$.

For completeness, we specify the one-period objective used in the long-run result. For a matching m and task vector v, let $Y_i(m, v) := \sum_k v_k m_{ik}$ be worker i's current task utility and let $\ell_i(m, v) := \sum_k \ell(v_k) m_{ik}$ be the smallest continuation promise that makes her assignment acceptable. For any vector h of participation floors, define $\Gamma_\lambda(h) := \max_{z \in \mathcal{U},\, z \geq h}\{\delta_P \Phi(z) + \delta_A \lambda \cdot z\}$, with value $-\infty$ when the constraint set is empty. After optimizing over continuation promises, the objective associated with current matching m is

$$W_\lambda(m; v) := (1 - \delta_P) \sum_{i,k} m_{ik} + (1 - \delta_A) \sum_i \lambda_i Y_i(m, v) + \Gamma_\lambda(\ell(m, v)). \tag{2}$$

For a task vector v, let $\mathcal{M}^{N,K}(v)$ be the matchings $m \in \mathcal{M}^{N,K}$ for which some $z \in \mathcal{U}$ satisfies $z \geq \ell(m, v)$. Thus $\mathcal{A}_\lambda(v)$ is the set of matchings in $\mathcal{M}^{N,K}(v)$ that maximize (2).

The proof uses costates and multipliers on (PC$_{ik}$) to study the assignments made by an arbitrary optimal contract. A costate $\lambda \in \Lambda(u)$ is a supporting slope of $-\Phi$ at the promise vector u. The multiplier μ_i measures the cost of raising worker i's minimum continuation promise. For any fixed optimal contract and supporting costate, Lagrangian optimality applies after almost every task realization. The next lemma attaches multipliers and next-period costates to the contract without changing its decisions.

A zero participation floor repeats the restriction $z_i \geq 0$ already imposed by the promise set, so we set its multiplier equal to zero. To choose among the possible multipliers on positive floors, let $I(\ell) := \{i : \ell_i > 0\}$. The proof shows that the realized positive floors lie in the interior of their feasible set with probability one. At such a floor vector ℓ, choose $\mu \geq 0$, with $\mu_i = 0$ outside $I(\ell)$, so that

$$\Gamma_\lambda(h) \leq \Gamma_\lambda(\ell) - \mu_{I(\ell)} \cdot (h_{I(\ell)} - \ell_{I(\ell)})$$

for every $h \in \mathcal{U}$ satisfying $h_i = 0$ outside $I(\ell)$. This inequality measures how the best continuation value changes when the positive floors change. If several vectors satisfy it, choose one with the smallest Euclidean norm. The proof shows that this choice exists, satisfies the costate equation for every selected optimal continuation, and can be made measurably.

Lemma 12 (Finite costates and multipliers for every optimal contract)**.** *Let an arbitrary optimal Markov contract start from $u_0 = 0$. On an event of probability one, there are finite vectors (λ_t, μ_t) at every finite date, with $\lambda_t \in \Lambda(u_t)$, $\lambda_0 = 0$, and $\mu_t \geq 0$, such that*

$$\lambda_{t+1} = q\lambda_t + \frac{\mu_t}{\delta_P}, \qquad \mu_{i,t}(u_{i,t+1} - \ell_{i,t}) = 0, \qquad \mu_{i,t} = 0 \text{ if } \ell_{i,t} = 0. \tag{3}$$

The costate λ_t is $\mathcal{F}_t$-measurable. Conditional on $\mathcal{F}_t$, μ_t and λ_{t+1} can be selected measurably as functions of the task vector and the contract's chosen matching. Thus the realized multiplier $\mu_t = \mu_t(V_t)$ is determined after the task draw. When the multiplier is not unique, it can be chosen by the minimum-norm rule above. These auxiliary choices leave the contract's matching and continuation rules unchanged.

Proof. Fix the current history and a finite current costate λ. Partition the task realizations according to the matching chosen by the contract and the set I of workers whose participation floors are positive. For a fixed matching and set I, let $\mathcal{D}_I := \{h_I \in \mathbb{R}^I_+ : (h_I, 0_{I^c}) \in \mathcal{U}\}$. Before conditioning on the contract's choice, the coordinates of ℓ_I are nonzero multiples of distinct task values, so ℓ_I has a joint density. Restricting this joint distribution to any event of positive probability preserves absolute continuity. The boundary of $\mathcal{D}_I$ is contained in the finite union of the hyperplanes $h_i = 0$ and $\sum_{i \in S} h_i = \overline{V}_{|S|}$ for every nonempty $S \subseteq I$. Hence ℓ_I lies in the interior of $\mathcal{D}_I$ almost surely. There are finitely many matchings and sets I, so this conclusion holds simultaneously across them.

For $h_I \in \mathcal{D}_I$, define $g_{\lambda,I}(h_I) := \Gamma_\lambda((h_I, 0_{I^c}))$. This is a finite concave function. At every interior point h_I, it has at least one finite supporting slope s_I, meaning that $g_{\lambda,I}(k_I) \leq g_{\lambda,I}(h_I) + s_I \cdot (k_I - h_I)$ for every $k_I \in \mathcal{D}_I$. The set of these slopes is nonempty, compact, and convex. Because raising a floor can only reduce the continuation value, every coordinate of s_I is nonpositive. Choose the slope of smallest Euclidean norm, set $\mu_I := -s_I$, and set $\mu_i := 0$ for $i \notin I$. When I is empty, set $\mu = 0$.

Let z be the continuation chosen by the contract. Lagrangian optimality implies that, after almost every task realization, z attains $\Gamma_\lambda((h_I, 0_{I^c}))$ for the selected matching. Recall that $A_\lambda(w) = \delta_P \Phi(w) + \delta_A \lambda \cdot w$. For every $w \in \mathcal{U}$, downward closure gives $(w_I, 0_{I^c}) \in \mathcal{U}$, and the supporting-slope inequality gives

$$A_\lambda(w) + \mu_I \cdot (w_I - h_I) \leq g_{\lambda,I}(w_I) + \mu_I \cdot (w_I - h_I) \leq g_{\lambda,I}(h_I) = A_\lambda(z).$$

Taking $w = z$, feasibility and $\mu \geq 0$ force equality, so $\mu_i(z_i - \ell_i) = 0$ for every i. The same inequality shows that z maximizes $A_\lambda(w) + \mu \cdot (w - \ell)$ over $w \in \mathcal{U}$. Therefore $\lambda^+ := q\lambda + \mu/\delta_P$ belongs to $\Lambda(z)$.

The smallest-norm slope is measurable in (λ, h_I). On the interior of $\mathcal{D}_I$, $g_{\lambda,I}(h_I)$ is jointly continuous in these variables. The supporting-slope inequalities are preserved under limits, and the slopes at each point form a compact set. Their unique smallest-norm element can therefore be selected measurably. The function $g_{\lambda,I}$ is differentiable almost everywhere on the interior of $\mathcal{D}_I$; wherever it is differentiable, the selected multiplier satisfies $\mu_I = -\nabla g_{\lambda,I}(h_I)$. This identity is used below to bound expected multipliers. Conditional on $\mathcal{F}_t$ and on any positive-probability cell of the finite partition by the selected matching and the set I of positive floors, the subvector ℓ_I is absolutely continuous by the argument above. It therefore avoids the exceptional set with probability one. When I is empty, there is no exceptional set to avoid. Define the multiplier and costate maps arbitrarily on the null sets where a floor lies on the boundary or the selected continuation fails the Lagrangian optimality condition. This gives measurable maps $v \mapsto (\mu(v), \lambda^+(v))$.

Finally, set $\lambda_0 = 0$, which belongs to $\Lambda(0)$ because zero lies on the plateau. Suppose $\lambda_t \in \Lambda(u_t)$ is finite and $\mathcal{F}_t$-measurable. Conditional on $\mathcal{F}_t$, Lagrangian optimality and the construction above produce measurable $\mu_t(v)$ and $\lambda_{t+1}(v) \in \Lambda(u_{t+1}(v))$ for almost every task vector, without changing the contract's matching or continuation. This proves the claim by induction. The exceptional event has conditional probability zero at each date, and a countable union over dates still has probability zero. Hence the stated equations hold at every finite date on one event of probability one. □

We will also use throughout the following direct implication of optimality: every positive-valued task is assigned after almost every draw. If a positive task were left unassigned, at most $K - 1$ tasks would be assigned. Since $K \leq N$, an idle worker could receive it with the continuation vector unchanged. This preserves (PC$_{ik}$), relaxes promise keeping, and strictly raises the principal's current payoff.

A.5.1 The case $q > 1$: eventual positive-only assignment

This subsection proves part (i) of Theorem 3 in two stages. First, we show that at least K workers acquire positive costates in finite time. Suppose that the set A of workers with positive costates stopped growing with $|A| < K$. Because $q > 1$, their costates would grow geometrically, forcing their total promise toward the feasible maximum for that group. At this boundary, the tasks allocated to A are determined, and the problem separates into a fixed contribution from A and a reduced dynamic problem for the remaining workers $Z = A^c$. The reduced problem cannot remain forever where promises to Z leave the principal's value unchanged. Leaving this set makes at least one costate in Z positive, and compactness makes the probability of doing so uniform near the boundary. Repeating this argument yields K positive costates in finite time.

The second stage shows that negative assignments eventually stop. Any negative task assigned after K costates are positive must lie in an interval below zero whose width is inversely proportional to the smallest positive costate. On each interval during which the set of positive costates is fixed, this minimum grows geometrically, and there are only finitely many such intervals. The conditional probabilities of a negative assignment therefore have a finite sum, so negative tasks are assigned at only finitely many dates.

Write $[N] := \{1, \dots, N\}$ and, for $S \subseteq [N]$, write $u(S) := \sum_{i \in S} u_i$. Recall that the promise set is described by the following subset-sum bounds:

$$\mathcal{U} = \{u \in \mathbb{R}_+^N : u(S) \le \overline{V}_{|S|} \text{ for every } S \subseteq [N]\},$$

where $\overline{V}_m := \overline{V}_K$ when $m > K$. Put $r_m := \overline{V}_m$, $r_0 := 0$, and $\beta_m := r_m - r_{m-1}$ for $m \le K$; thus β_m is the increase in the maximum deliverable utility when the group size rises from $m - 1$ to m. These increments satisfy $\beta_1 \ge \cdots \ge \beta_K > 0$. Let $p_a := (r_K - r_a)/(K - a)$ for $a < K$. The maintained condition in the theorem is $\mathbb{E}[V] < \bar{h}$.

Fix an arbitrary optimal Markov contract. By Lemma 12, it can be paired, without changing its matching or continuation-promise rules, with finite supporting costates and multipliers on (PC_{ik}) at every finite date, satisfying (3). Call a worker *active* when the worker's costate is strictly positive. Once a worker becomes active, the recursion keeps that worker active. Throughout this subsection, costates and multipliers are selected as in Lemma 12, so they satisfy (3) without changing the contract's decisions.

Two preliminary bounds. We begin with two bounds used later when fewer than K costates are positive: one on feasible promises and one on the values of the remaining tasks. For any $A \subseteq [N]$, with $a := |A|$, define

$$\mathcal{U}_A := \{x_A \geq 0 : x_A(R) \leq r_{|R|} \text{ for every } R \subseteq A\}, \qquad \mathcal{B}_A := \{x_A \in \mathcal{U}_A : x_A(A) = r_a\}.$$

When $a < K$, the total promise available after A reaches its maximum is $r_K - r_a = (K-a)p_a$.

Lemma 13 (Completing and extending feasible promise vectors)**.** *(i) For every $x_A \in \mathcal{U}_A$, there is a $b_A \in \mathcal{B}_A$ such that $b_A \geq x_A$ coordinatewise.*

(ii) Fix $x_A \in \mathcal{U}_A$ and put $Z := A^c$. The promises for Z that complete x_A to a vector in $\mathcal{U}$ form the set

$$\mathcal{C}_Z(x_A) = \left\{ z_Z \geq 0 : z_Z(T) \leq \min_{B \subseteq A}\{r_{|B|+|T|} - x_A(B)\} \text{ for every } T \subseteq Z \right\}.$$

If $a < K$, then $p_a \mathbf{1}_S \in \mathcal{C}_Z(x_A)$ for every $S \subseteq Z$ with $|S| = K - a$.

Proof. For part (i), maximize $b_A(A)$ over the nonempty compact set of $b_A \in \mathcal{U}_A$ satisfying $b_A \geq x_A$. Suppose a maximizer has $b_A(A) < r_a$. The maximizer is coordinatewise maximal, so every $i \in A$ belongs to a set $S_i \subseteq A$ satisfying $b_A(S_i) = r_{|S_i|}$. If the constraints for S and T bind, then

$$\begin{aligned} r_{|S|} + r_{|T|} &= b_A(S \cup T) + b_A(S \cap T) \\ &\leq r_{|S \cup T|} + r_{|S \cap T|} \leq r_{|S|} + r_{|T|}. \end{aligned}$$

The first inequality follows from feasibility, and the second from submodularity of r_m. Hence the constraint for $S \cup T$ also binds. Iterating over the sets S_i makes the constraint for A bind, a contradiction. Thus $b_A(A) = r_a$.

For part (ii), apply the constraint for $B \cup T$, where $B \subseteq A$ and $T \subseteq Z$, and then minimize over B. This gives the displayed bounds for $\mathcal{C}_Z(x_A)$. Conversely, these bounds imply the constraint for every $B \cup T$, and every subset of $A \cup Z$ has this form. They therefore characterize $\mathcal{C}_Z(x_A)$.

Now suppose $a < K$ and fix $S \subseteq Z$ with $|S| = K - a$. For $B \subseteq A$ and a nonempty $T' \subseteq S$, write $b := |B|$ and $t := |T'|$. Because the increments β_m are decreasing,

$$\frac{r_{b+t} - r_b}{t} = \frac{1}{t}\sum_{j=1}^{t} \beta_{b+j} \geq \frac{1}{t}\sum_{j=1}^{t} \beta_{a+j} \geq \frac{1}{K-a}\sum_{j=1}^{K-a} \beta_{a+j} = p_a.$$

For arbitrary $T \subseteq Z$, let $T' := T \cap S$. If $T' \neq \varnothing$, feasibility of x_A and the preceding inequality give $x_A(B) + p_a|T'| \leq r_{b+|T'|} \leq r_{b+|T|}$. If $T' = \varnothing$, the same conclusion follows directly from $x_A(B) \leq r_b$. Thus $(x_A, p_a \mathbf{1}_S)$ satisfies every constraint defining $\mathcal{U}$. □

We next bound the task values left for workers outside A. This bound supplies the strict inequality used below to show that their value cannot remain at its maximum forever.

Lemma 14 (Average value of the remaining tasks)**.** *Let $M \geq 1$ and $0 \leq a < M$ be integers. Let $X_1, \ldots, X_M$ be i.i.d. with mean μ, $\Pr(X_1 > 0) > 0$, and $\Pr(X_1 < 0) > 0$. Perform a steps, and in each step, remove the highest-valued positive task still available; if every remaining task is nonpositive, leave the set of tasks unchanged. After these steps, let T_{M-a} be the sum of the $M - a$ highest-valued tasks still available. If $\mathbb{E}[T_{M-a}] \geq 0$, then $\mathbb{E}[T_{M-a}] \leq (M - a)\mu$.*

Proof. Write $L^0 := T_{M-a}/(M-a)$. Represent each task as $X_i = Q(U_i)$, where the U_i are independent and uniform on $[0, 1]$ and Q is a quantile function. When two task values tie, rank the task with the higher U_i higher. This tie-breaking rule does not change T_{M-a}. Let I_i indicate that task i is among the $M - a$ tasks counted in T_{M-a}. Exchangeability gives

$$\mathbb{E}[L^0] = \int_0^1 Q(u)\alpha(u)\, du, \qquad \alpha(u) := \frac{M}{M-a} \Pr(I_1 = 1 \mid U_1 = u).$$

Let $p_+ := \Pr(X > 0)$, $p_- := \Pr(X < 0)$, and $H(t) := \Pr\{\mathrm{Bin}(M-1, t) \geq a\}$. If $Q(u) > 0$, task 1 is counted precisely when at least a other uniforms exceed u. If $Q(u) < 0$, it is counted precisely when at least a other uniforms are either in the positive tail or below u. Therefore

$$\alpha(u) = \frac{M}{M-a} \begin{cases} H(p_+ + u), & Q(u) < 0, \\ H(1-u), & Q(u) > 0. \end{cases}$$

The value of α where $Q(u) = 0$ is immaterial. The function H is increasing. Moreover, integrating over an additional independent uniform shows that $\int_0^1 H(t)\, dt = (M-a)/M$: the rank of that uniform among M uniforms is itself uniform.

The integral of α over a positive upper interval of length $x \leq p_+$ and over a negative lower interval of length $y \leq p_-$ is, respectively,

$$A_+(x) := \frac{M}{M-a} \int_0^x H(t)\, dt, \qquad A_-(y) := \frac{M}{M-a} \int_{p_+}^{p_+ + y} H(t)\, dt,$$

and $B_+(x) := x - A_+(x)$ and $B_-(y) := y - A_-(y)$. Monotonicity of H and its integral identity imply $A_+(x)/x \leq 1$ and $A_+(x)/x \leq A_-(y)/y$ whenever $x, y > 0$. Hence

$$\frac{B_+(x)}{A_+(x)} \geq \frac{B_-(y)}{A_-(y)}$$

for $x, y > 0$. Choose

$$c := \max\left\{0, \sup_{0<y\le p_-} \frac{B_-(y)}{A_-(y)}\right\}.$$

Then $B_+(x) \ge cA_+(x)$ and $B_-(y) \le cA_-(y)$. Let $Q_+ := \max\{Q, 0\}$ and $Q_- := \max\{-Q, 0\}$, and define $P_w := \int_0^1 Q_+(u)w(u)\,du$ and $N_w := \int_0^1 Q_-(u)w(u)\,du$. For $t > 0$, let x_t be the length of the upper interval $\{u : Q(u) > t\}$ and y_t the length of the lower interval $\{u : Q(u) < -t\}$. The identities $Q_+(u) = \int_0^\infty \mathbf{1}\{Q(u) > t\}\,dt$ and $Q_-(u) = \int_0^\infty \mathbf{1}\{Q(u) < -t\}\,dt$, together with Fubini's theorem, give

$$\begin{aligned} P_\alpha &= \int_0^\infty A_+(x_t)\,dt, \quad P_{1-\alpha} = \int_0^\infty B_+(x_t)\,dt, \\ N_\alpha &= \int_0^\infty A_-(y_t)\,dt, \quad N_{1-\alpha} = \int_0^\infty B_-(y_t)\,dt. \end{aligned}$$

Consequently, $P_{1-\alpha} \ge cP_\alpha$ and $N_{1-\alpha} \le cN_\alpha$. Thus $P_\alpha - N_\alpha \ge 0$ implies $P_{1-\alpha} - N_{1-\alpha} \ge c(P_\alpha - N_\alpha) \ge 0$. Equivalently,

$$\int_0^1 Q(u)\alpha(u)\,du \ge 0 \quad \Longrightarrow \quad \int_0^1 Q(u)(1-\alpha(u))\,du \ge 0.$$

The second integral equals $\mu - \mathbb{E}[L^0]$, proving the claim. □

Linking constant-value regions to positive costates. For $Z \subseteq [N]$ and $A := Z^c$, define

$$\mathcal{G}_Z := \{u \in \mathcal{U} : \Phi(u) = \Phi(u_A, 0_Z)\}.$$

At these states, setting the promises of the workers in Z to zero leaves the principal's value unchanged.

Lemma 15 (Zero costates and changes in value)**.** *The set $\mathcal{G}_Z$ is compact. If $\lambda \in \Lambda(u)$ and $\lambda_Z = 0$, then $u \in \mathcal{G}_Z$. On the probability-one event on which (3) holds, if $\lambda_Z = 0$ and a selected next promise u^+ lies outside $\mathcal{G}_Z$, then $\lambda_j^+ > 0$ for some $j \in Z$.*

Proof. The set $\mathcal{G}_Z$ is the zero set of the continuous function $u \mapsto \Phi(u) - \Phi(u_A, 0_Z)$ and is therefore a closed subset of the compact set $\mathcal{U}$. Monotonicity gives $-\Phi(u_A, 0_Z) \le -\Phi(u)$. The subgradient inequality for $-\Phi$, evaluated at $(u_A, 0_Z)$, gives the reverse inequality when $\lambda_Z = 0$.

For the last statement, let μ be the multiplier on (PC_{ik}) in the costate recursion and suppose $\mu_Z = 0$. The recursion then gives $\lambda_Z^+ = 0$. Since λ^+ supports $-\Phi$ at u^+, the preceding conclusion would imply $u^+ \in \mathcal{G}_Z$. Thus $u^+ \notin \mathcal{G}_Z$ requires $\mu_j > 0$ for some $j \in Z$. Because $\lambda_j = 0$, the recursion gives $\lambda_j^+ = \mu_j/\delta_P > 0$. □

The next compactness result will be used to obtain a uniform probability of leaving a constant-value set and to extend that probability to nearby states. A *future-history law* from promise u is the joint distribution of all subsequent task draws, matchings, and continuation promises under a feasible contract, including its public randomization.

Lemma 16 (Compact sets of decisions and future-history laws)**.** *The pairs (u, Q) for which Q is a feasible randomized one-period decision at u form a compact set under weak convergence; the Bellman-optimal pairs form a compact subset. The pairs (u, P) for which P is a feasible future-history law from u also form a compact set, as do the optimal pairs. Each decision or law can be chosen Borel measurably as a function of u.*

Proof. The compactness argument in Appendix A.1 gives the first set of one-period pairs. The Bellman payoff is continuous, so imposing equality with the continuous value function Φ gives a closed, hence compact, set of Bellman-optimal pairs.

For future-history laws, let

$$\Omega := \mathcal{U} \times \left([\underline{v}, \overline{v}]^K \times \mathcal{M}^{N,K} \times \mathcal{U}\right)^{\mathbb{N}}.$$

Give $\mathcal{M}^{N,K}$ the discrete topology. Then Ω is compact and metrizable, and the space of probability laws on Ω is compact under weak convergence. Write its canonical coordinates as $(u_0, (v_t, m_t, u_{t+1})_{t\geq 0})$, and let $h_t := (u_0, v_0, m_0, u_1, \ldots, v_{t-1}, m_{t-1}, u_t)$ denote the pre-draw history through u_t.

For a law starting from promise u, feasibility requires $u_0 = u$, $u_{i,t+1} \geq \ell_i(m_t, v_t)$ for every i and t, conditional task law F^K, and conditional promise keeping. For every nonnegative continuous function g of h_t, promise keeping is equivalent to

$$\int g(h_t)\{(1 - \delta_A)Y_i(m_t, v_t) + \delta_A u_{i,t+1} - u_{i,t}\}\, dP \geq 0$$

for every i. For every continuous g of h_t and every continuous k, the conditional task-law restriction is

$$\int g(h_t)k(v_t)\, dP = \int g(h_t)\, dP \int k(v)\, dF^K(v)$$

at every date. The matching restriction is built into the definition of Ω. The paths satisfying (PC_{ik}) form a closed subset of Ω, and the laws that assign this set probability one form a closed set under weak convergence. The two displayed restrictions are closed because their integrands are bounded and continuous. The initial-state condition is also closed under joint convergence of (u, P). Hence the set of feasible pairs (u, P) is compact.

Let $J(P)$ denote the normalized discounted number of assignments under P. Its truncation after date $L - 1$ is a continuous function of P, and the omitted tail is bounded uniformly by $K\delta_P^L$. Thus J is continuous. Self-generation gives at least one feasible law at every promise, and compactness implies that an optimal law exists. The optimal pairs (u, P) are exactly the feasible pairs satisfying $J(P) = \Phi(u)$, so they form a compact set. For each promise, all four sets of available choices are nonempty and compact, and each set of pairs is Borel. The standard measurable-selection theorem gives the stated choices. □

Assignments at the maximal-promise boundary. We now characterize the boundary approached when the set of positive costates remains fixed. For a group A of size a, the equality $u(A) = r_a$ means that the group has reached the largest total promise it can feasibly receive. The following result shows that this equality persists and determines the tasks assigned to A.

Lemma 17 (Behavior when a group's total promise is maximal)**.** *Let $A \subseteq [N]$ and $a := |A|$. If $u(A) = r_a$, every feasible one-period policy from u has $z(A) = r_a$ almost surely. The workers in A receive every positive task among the* $\min\{a, K\}$ *highest-valued tasks and receive no negative task, almost surely. Consequently, every future promise remains in $\{w \in \mathcal{U} : w(A) = r_a\}$.*

Proof. Let $S_a(V)$ be the sum of the positive values among the $\min\{a, K\}$ highest-valued tasks, and let Y_A be the total current task value assigned to the workers in A. Pointwise, $Y_A \le S_a(V)$, and $\mathbb{E}[S_a(V)] = r_a$. Feasibility also gives $z(A) \le r_a$ pointwise. Summing promise keeping over A yields

$$r_a \le (1 - \delta_A)\mathbb{E}[Y_A] + \delta_A \mathbb{E}[z(A)] \le (1 - \delta_A) r_a + \delta_A r_a = r_a.$$

Thus $\mathbb{E}[S_a(V) - Y_A] = 0$ and $\mathbb{E}[r_a - z(A)] = 0$. Both variables are nonnegative, so both equal zero almost surely. The density assumption rules out ties and zero-valued tasks almost surely, giving the stated assignment. Applying the same argument at each continuation proves the final claim. □

Reducing the problem to the remaining workers. Once $u(A) = r_a$, the preceding lemma fixes the contribution of A. We can therefore separate A from $Z = A^c$ and study the reduced problem for Z using the tasks left after A has been served. Fix $A \subseteq [N]$ with $a := |A| < K$, put $Z := A^c$, and let $\mathcal{F}_A := \{u \in \mathcal{U} : u(A) = r_a\}$. Define the promise set left for Z by

$$\mathcal{C}_Z^a := \{z_Z \ge 0 : z_Z(B) \le r_{a+|B|} - r_a \text{ for every } B \subseteq Z\}.$$

On $\mathcal{F}_A$, Lemma 17 requires A's matching to attain the sum of the positive values among the a highest-valued tasks. Represent each action by disjoint matchings for A and Z, and require the matching for A to attain this sum. Let Ψ_a be the Bellman value that counts the assignments to Z, whose promise set is $\mathcal{C}_Z^a$. The decision rule observes the full task vector and all public randomization. At ties, allow every maximizing matching for A; this relation is closed.

The arguments of Lemmas 2 and 16 apply to this reduced problem. Its state space is the compact polytope above, its joint matching set is finite, and its feasibility and Bellman restrictions are closed. Thus its feasible one-period decisions and its feasible and optimal future-history laws form compact sets of pairs and admit Borel measurable choices.

Lemma 18 (Separating the promises of the two worker groups)**.** *The set $\mathcal{F}_A$ factors as $\mathcal{F}_A = \mathcal{B}_A \times \mathcal{C}_Z^a$. There is a constant C_a such that*

$$\Phi(x_A, z_Z) = C_a + \Psi_a(z_Z) \qquad \textit{for every } (x_A, z_Z) \in \mathcal{F}_A.$$

The value Ψ_a for the remaining workers is continuous, concave, symmetric, and coordinatewise nonincreasing. If $\mathcal{P}_a := \{z_Z : \Psi_a(z_Z) = \Psi_a(0)\}$, then

$$\mathcal{G}_Z \cap \mathcal{F}_A = \mathcal{B}_A \times \mathcal{P}_a.$$

Proof. If $(x_A, z_Z) \in \mathcal{F}_A$, the constraint for $A \cup B$ gives $z_Z(B) \le r_{a+|B|} - r_a$. Conversely, take $x_A \in \mathcal{B}_A$ and $z_Z \in \mathcal{C}_Z^a$. For $R \subseteq A$ and $B \subseteq Z$, the decreasing increments of r_m, including its zero increments after K, give

$$r_{|R|} + r_{a+|B|} - r_a \le r_{|R|+|B|}.$$

Therefore $x_A(R) + z_Z(B) \le r_{|R|+|B|}$, proving the product representation.

By Lemma 17, A receives every positive task among the a highest-valued tasks, and every continuation remains in $\mathcal{F}_A$. The reduced action records a maximizing matching for A. Outside task ties, every such matching removes the same task indices: the positive tasks among the a highest-valued tasks. Thus the worker-order lottery used below to implement x_A leaves the same task set available to Z, and the choices for the two groups can be separated after the task draw. We next verify that $\mathcal{C}_Z^a$ is exactly the promise set available to the remaining workers. The set $\mathcal{B}_A$ is the convex hull of the coordinate permutations of $(\beta_1, \dots, \beta_a)$. At each date, draw a fresh public lottery over the orders in which the workers in A select positive tasks. In each order they receive every positive task among the a highest-valued tasks; varying the probabilities over orders delivers any prescribed $x_A \in \mathcal{B}_A$ in expectation. Holding their continuation equal to x_A makes each of their promise-keeping inequalities bind.

After A has been served, the extreme points of $\mathcal{C}_Z^a$ with maximal total promise are the permutations in $\mathbb{R}^Z$ of $(\beta_{a+1}, \dots, \beta_K, 0, \dots, 0)$. Every $z_Z \in \mathcal{C}_Z^a$ is dominated by some point with maximal total promise $\widetilde{z}_Z$, by the argument in part (i) of Lemma 13. At each date, publicly randomize over worker orders that give expected current utility $\tilde{z}_Z$, and hold the continuation at z_Z. The expected delivered utility is $(1-\delta_A)\tilde{z}_Z + \delta_A z_Z \geq z_Z$. Constraint ($\mathrm{PC}_{ik}$) holds because only positive tasks are assigned. Hence this rule delivers every promise in $\mathcal{C}_Z^a$ and returns the continuation to that set. Necessity follows from the $A \cup B$ subset constraints. Thus $\mathcal{C}_Z^a$ is exactly the promise set for Z.

Let $n_A(v)$ be the number of positive tasks received by A. This number equals $\min\{a, \#\{k : v_k > 0\}\}$ independently of the choices for Z. Project any feasible future-history law starting in $\mathcal{F}_A$ by retaining the complete task vector and the outcomes of all public randomization, removing A's actions, and keeping the matching and continuation coordinates of Z. The resulting future-history law for Z is feasible, and the full payoff is its payoff plus $n_A(V_t)$ at every date. Since task draws have the same distribution at every date, the normalized discounted contribution from A is $\mathbb{E}[n_A(V)]$. Thus such a law has value at most $\mathbb{E}[n_A(V)] + \Psi_a(z_Z)$.

For the reverse construction, start with an optimal future-history law for Z. At every date, draw the public worker-order lottery that implements x_A, give A every positive task among the a highest-valued tasks, keep its continuation at x_A, and apply the law for Z to the task indices left over. The law for Z may condition on the full task vector and all public randomization. The two matchings are disjoint, promise keeping for A binds, and the product representation of the promise set makes the combined continuation feasible. This extends the law for Z to the full problem on $\mathcal{F}_A$ and attains $\mathbb{E}[n_A(V)] + \Psi_a(z_Z)$. Hence $C_a = \mathbb{E}[n_A(V)]$ and the value decomposition follows.

Mixing future-history laws for Z gives concavity, relabeling the workers gives symmetry, and a law feasible at one promise remains feasible at every lower promise, giving coordinatewise monotonicity. The continuity argument in Lemma 2 applies to this reduced problem. Thus $\mathcal{P}_a$ is nonempty, compact, symmetric, and downward closed. The last displayed identity follows directly from the value decomposition. □

The definitions also cover $A = \varnothing$. In that case $\mathcal{F}_A = \mathcal{U}$, $\mathcal{C}_Z^0 = \mathcal{U}$, $\Psi_0 = \Phi$, and $\mathcal{P}_0 = \mathcal{P}$.

For a task draw, remove every positive task among the a highest-valued tasks, and let T_s^a be the sum of the $s := K - a$ highest-valued tasks that remain. Applying Lemma 14 with $M = K$, $X_i = V_i$, $a = |A|$, and the maintained inequality $\mathbb{E}[V] < \bar{h}$ gives

$$\mathbb{E}[T_s^a] < s\bar{h}. \tag{4}$$

Indeed, a negative left side satisfies the claim immediately; if it is nonnegative, Lemma 14 bounds it above by $s\mathbb{E}[V] < s\bar{h}$. The expected sum of all positive tasks left after A has been served is $r_K - r_a = sp_a$.

Why the remaining workers must leave the constant-value set. The set $\mathcal{P}_a$ contains exactly the promises for Z at which their value Ψ_a equals its value at zero. The next lemma uses the preliminary promise and task-value bounds to show that remaining in this set forever is strictly suboptimal. Compactness then yields a uniform positive probability of leaving within a fixed number of dates.

Lemma 19 (Leaving the remaining workers' constant-value set)**.** *Every optimal future-history law for Z starting from a state in $\mathcal{P}_a$ leaves $\mathcal{P}_a$ in finite time almost surely. There are also constants $L_a < \infty$ and $c_a > 0$ such that every such law starting at any state in $\mathcal{P}_a$ leaves within L_a dates with probability at least c_a.*

Proof. Let $V^{\mathcal{P}_a}(0)$ be the value at zero when all future promises are constrained to $\mathcal{P}_a$. If an optimal law from any $z \in \mathcal{P}_a$ remained in $\mathcal{P}_a$ forever almost surely, it would be feasible from zero and would have value $\Psi_a(z) = \Psi_a(0)$. It therefore suffices to prove $V^{\mathcal{P}_a}(0) < \Psi_a(0)$.

The largest common floor feasible for any s coordinates of Z is p_a. Indeed, the relevant constraints are $th \leq r_{a+t} - r_a$ for $1 \leq t \leq s$, and the decreasing differences β_m make the smallest ratio occur at $t = s$.

Define

$$\rho_a := \max\{h \geq 0 : h\mathbf{1}_S \in \mathcal{P}_a \text{ for some } S \subseteq Z,\ |S| = s\}.$$

The maximum exists because $\mathcal{P}_a$ is compact. Symmetry makes the choice of S immaterial, downward closure includes all $h \leq \rho_a$, and feasibility gives $\rho_a \leq p_a$.

First suppose $\rho_a < \min\{p_a, \bar{h}\}$. Choose $\eta > 0$ such that $\rho_a + \eta < \min\{p_a, \bar{h}\}$ and $\delta_P \omega_a(\eta) < 1 - \delta_P$, where ω_a is a modulus of continuity for Ψ_a. Because F has a strictly positive density, the event on which all K participation floors belong to $(\rho_a, \rho_a + \eta)$ has positive probability. Every task is negative on this event, so the workers in A receive no task. A continuation in $\mathcal{P}_a$ can satisfy (PC_{ik}) for at most $s - 1$ assignments: s assignments would imply, by (PC_{ik}) and downward closure, that $h\mathbf{1}_S \in \mathcal{P}_a$ for some $h > \rho_a$.

Starting from zero, take an optimal policy whose continuations are constrained to $\mathcal{P}_a$. On this event, replace its decision by assigning s remaining tasks to distinct workers, setting those workers' continuation promises equal to their participation floors, and setting every other continuation promise to zero. This vector is feasible by part (ii) of Lemma 13, because all its coordinates are below p_a, and it is within sup-norm η of $\rho_a \mathbf{1}_S \in \mathcal{P}_a$. Each assigned negative task together with its floor delivers zero utility, so promise keeping at zero is preserved. Conditional on the event, the principal gains at least one current assignment and loses at most $\omega_a(\eta)$ in continuation value. The gain in normalized payoff is strictly positive because $1 - \delta_P - \delta_P \omega_a(\eta) > 0$. Hence $\Psi_a(0) > V^{\mathcal{P}_a}(0)$.

It remains to cover the cases in which $\mathcal{P}_a$ reaches one of the two relevant boundaries. Suppose, toward a contradiction, that $\Psi_a(0) = V^{\mathcal{P}_a}(0)$. Attainment and Lemma 16 make the following set nonempty and compact:

$$\mathfrak{H}_a := \{(z, P) : z \in \mathcal{P}_a, P \text{ is optimal from } z, P(z_t \in \mathcal{P}_a \text{ for every } t) = 1\}.$$

The principle of optimality implies that, after almost every first-period history, the next state and conditional tail law also belong to $\mathfrak{H}_a$. Let

$$M_a := \max\{z(Z) : (z, P) \in \mathfrak{H}_a\}.$$

At a maximizing pair, every next promise for Z has total at most M_a. Summing promise keeping and using this upper bound on every next total gives

$$M_a \leq (1 - \delta_A)\mathbb{E}[Y_Z] + \delta_A M_a, \qquad \text{and hence} \qquad M_a \leq \mathbb{E}[Y_Z], \tag{5}$$

where Y_Z is the total value of the tasks currently assigned to workers in Z.

A policy feasible from a promise in $\mathcal{P}_a$ is feasible from zero because promise keeping is a lower bound. Since Ψ_a is constant on $\mathcal{P}_a$, every law in $\mathfrak{H}_a$ is therefore optimal from zero. At zero, every decision satisfying matching feasibility and (PC$_{ik}$) pointwise also satisfies promise keeping. An alternative that makes strictly more assignments and chooses its continuation in $\mathcal{P}_a$ has the same continuation value and a strictly higher current payoff. This comparison applies separately to each task realization.

Consider first $p_a < \bar{h}$ and $\rho_a = p_a$. For every $\varepsilon \in (0, p_a)$, there is a positive-probability event on which exactly a tasks are positive and the other s tasks have participation floors in $(p_a - \varepsilon, p_a)$. The workers in A take all positive tasks, and the vector of the other s floors belongs to $\mathcal{P}_a$ by downward closure. The preceding pointwise comparison implies that every law in $\mathfrak{H}_a$ assigns all s negative tasks. Each resulting next promise has total at least $s(p_a - \varepsilon)$, so $M_a \geq s(p_a - \varepsilon)$. Letting $\varepsilon \downarrow 0$ gives $M_a \geq sp_a = r_K - r_a$. Yet Y_Z is pointwise no larger than the sum of all positive tasks left after A has been served. Its expectation is $r_K - r_a$, and the inequality is strict on the event just described because $Y_Z < 0$ there. Thus $M_a \leq \mathbb{E}[Y_Z] < r_K - r_a$, contradicting the lower bound on M_a.

Now consider $\bar{h} \leq p_a$ and $\rho_a \geq \bar{h}$. The participation floors are strictly below $\bar{h}$ almost surely. The same pointwise comparison implies that every law in $\mathfrak{H}_a$ assigns at least s tasks to Z after almost every draw: if it assigned fewer, it could assign s available tasks and use their floor vector, which is dominated by $\bar{h}\mathbf{1}_S \in \mathcal{P}_a$. For every $\varepsilon \in (0, \bar{h})$, the event on which all K participation floors belong to $(\bar{h} - \varepsilon, \bar{h})$ has positive probability. On this event the next promise to Z has total at least $s(\bar{h} - \varepsilon)$. Hence $M_a \geq s(\bar{h} - \varepsilon)$, and letting $\varepsilon \downarrow 0$ gives $M_a \geq s\bar{h}$. If $d \geq s$ tasks are assigned to Z, their sum is at most the sum of the d highest-valued tasks left after serving A, which is at most T_s^a because all terms after rank s among those remaining tasks are nonpositive.

Using (5) and (4) gives

$$s\bar{h} \leq M_a \leq \mathbb{E}[Y_Z] \leq \mathbb{E}[T_s^a] < s\bar{h},$$

a contradiction. These cases cover every possibility because $\rho_a \leq p_a$. We have proved $\Psi_a(0) > V^{\mathcal{P}_a}(0)$ in all cases. By the reduction at the start of the proof, no optimal law from a state in $\mathcal{P}_a$ can remain there with probability one.

We next make the probability uniform. If no pair (L_a, c_a) existed, for each n there would be a state $z_n \in \mathcal{P}_a$ and an optimal law P_n with probability below $1/n$ of leaving by date n. Compactness gives a subsequence converging to an optimal law P_* from some $z_* \in \mathcal{P}_a$. For every fixed L, the event that the first L states belong to the closed set $\mathcal{P}_a$ is closed and has P_n-probability tending to one. The Portmanteau theorem gives it P_*-probability one. Taking the countable intersection over L produces an optimal law that stays on $\mathcal{P}_a$ forever, a contradiction. A uniform pair (L_a, c_a) therefore exists. After any history at which a law remains in $\mathcal{P}_a$, the principle of optimality makes its conditional continuation an optimal law from the current promise. Applying the uniform bound successively over blocks of L_a dates gives probability at most $(1 - c_a)^n$ of remaining for n blocks. Thus every optimal law leaves $\mathcal{P}_a$ in finite time almost surely. □

From the boundary to another positive costate. Along a path on which the set of positive costates remains A, the total promise to A approaches r_a. At the boundary, leaving $\mathcal{P}_a$ in the reduced problem is equivalent to leaving $\mathcal{G}_Z$ in the full problem. As long as $\lambda_Z = 0$, such an exit makes a costate in Z positive. The compactness result extends the uniform probability of this event to nearby states.

Lemma 20 (A zero costate becomes positive near the promise boundary)**.** *Fix a nonempty $A \subsetneq [N]$ with $a := |A| < K$ and put $Z := A^c$. There are constants $\eta_A > 0$, $L_A < \infty$, and $d_A > 0$ such that the following holds. If $r_a - u(A) \leq \eta_A$ and a finite supporting costate satisfies $\lambda_i > 0$ for $i \in A$ and $\lambda_Z = 0$, then, under every optimal continuation with costates satisfying* (3)*, some costate coordinate in Z becomes positive within L_A dates with probability at least d_A.*

Proof. Take L_A from Lemma 19. Suppose the conclusion failed. For every n, there would be a state u^n with $r_a - u^n(A) \leq 1/n$, a finite supporting costate λ^n that is positive on A and zero on Z, and an optimal continuation law under which the probability that a costate in Z becomes positive during the next L_A dates is below $1/n$. Compactness of the optimal pairs in Lemma 16 gives, after passing to a subsequence, $u^n \to u^* \in \mathcal{F}_A$ and weak convergence to an optimal continuation law from u^*.

By Lemma 15, every u^n belongs to $\mathcal{G}_Z$, and every selected state before a costate in Z becomes positive also belongs to this closed set. Therefore

$$\Pr(\text{leave } \mathcal{G}_Z \text{ by } L_A) \le \Pr(\text{a costate in } Z \text{ becomes positive by } L_A).$$

For a fixed finite horizon, the event that all continuation states belong to $\mathcal{G}_Z$ is closed. The Portmanteau theorem implies that the limiting law stays in $\mathcal{G}_Z$ for all L_A dates. Since its initial state is on $\mathcal{F}_A$, Lemma 17 keeps the limit law in $\mathcal{F}_A$. Its projection onto the workers in Z is a feasible future-history law from the projected initial state u^*_Z, and it is optimal there: replacing it by a strictly better future-history law for Z and applying the extension construction in Lemma 18 would strictly improve the full contract on $\mathcal{F}_A$. The identity $\mathcal{G}_Z \cap \mathcal{F}_A = \mathcal{B}_A \times \mathcal{P}_a$ therefore produces an optimal law for Z that remains in $\mathcal{P}_a$ throughout the next L_A dates. This contradicts Lemma 19, which gives probability at least c_a of leaving within L_A dates. The contradiction proves the existence of $\eta_A > 0$ and $d_A > 0$. □

We can now iterate this argument. Leaving the original plateau produces the first positive costate. Thereafter, if fewer than K costates are positive and their set stops growing, the current group approaches the boundary where the preceding lemma gives a probability bounded away from zero that another costate becomes positive.

Proposition 11 (At least K costates become positive in finite time)**.** *Suppose $q > 1$, $\mathbb{E}[V] < \ell(\underline{v})$, and $u_0 = 0$. Along every costate process supplied by Lemma 12, at least K costate coordinates become strictly positive in finite time almost surely.*

Proof. Apply Lemma 19 with $A = \varnothing$. In this case $\Psi_0 = \Phi$ and $\mathcal{P}_0 = \mathcal{P}$, so every optimal law starting at 0 leaves $\mathcal{P}$ in finite time. If a costate becomes positive before then, one positive coordinate already exists. Otherwise the selected costate remains zero until the transition that leaves $\mathcal{P}$. On that transition, Lemma 15, with $Z = [N]$, implies that at least one coordinate of the next costate is positive.

Once a costate becomes positive, it remains positive. Suppose, on an event of positive probability, that fewer than K coordinates ever become positive. The active set can expand only finitely many times. Since there are finitely many sets and countably many dates, it is enough to fix a nonempty A with $|A| < K$ and a deterministic date T, and rule out the event that the active set equals A at every date from T onward. On this event, the total promise to A approaches its maximum r_a. To see this, put $a := |A|$, $f := -\Phi$, and $\Omega_f := \max_{\mathcal{U}} f - \min_{\mathcal{U}} f$. At each date, use part (i) of Lemma 13 to choose $b_A \in \mathcal{B}_A$ with $b_A \ge u_{t,A}$, and set $b_Z = 0$. The resulting vector b belongs to $\mathcal{U}$. Since the active set has not changed, $\lambda_{t,Z} = 0$. The subgradient inequality gives

$$\min_{i \in A} \lambda_{i,t}\{r_a - u_t(A)\} \leq \lambda_t \cdot (b - u_t) \leq f(b) - f(u_t) \leq \Omega_f.$$

The costate recursion gives $\min_{i \in A} \lambda_{i,t} \geq q^{t-T} \min_{i \in A} \lambda_{i,T}$. Hence $r_a - u_t(A) \leq \Omega_f / (q^{t-T} \min_{i \in A} \lambda_{i,T}) \to 0$.

Define the pre-draw stopping time

$$\tau_A := \inf \left\{ t \geq T : \frac{\Omega_f}{q^{t-T} \min_{i \in A} \lambda_{i,T}} \leq \eta_A \text{ and the active set is } A \right\}.$$

It is finite on the event under consideration. Starting at τ_A, Lemma 20 gives conditional probability at least d_A that a costate outside A becomes positive during the next L_A dates. Conditional on failure, the active set is still A, and the preceding bound on the distance from $u(A) = r_a$ remains within the same neighborhood at every later block boundary. Iterating at $\tau_A + nL_A$ bounds the conditional probability of no such change through n blocks by $(1 - d_A)^n$. Letting $n \to \infty$ makes the probability of persistence zero. Taking the finite union over A and the countable union over T proves that the active set reaches size K in finite time almost surely. □

After K costates are positive: negative assignments eventually stop. It remains to complete the second stage of the argument. Once K costates are positive, the corresponding workers can receive a continuation vector with total r_K, the largest feasible total promise. The next lemma uses this fact to bound every negative task assigned by an optimal policy, including one assigned to a worker with zero costate.

Lemma 21 (Bound on the magnitude of an assigned negative task)**.** *Fix a state u and a supporting costate $\lambda \in \Lambda(u)$. Suppose $A(\lambda) := \{i : \lambda_i > 0\}$ has at least K members, and put $a(\lambda) := \min_{i \in A(\lambda)} \lambda_i > 0$. For any fixed optimal policy, after almost every task draw, every assigned value $x < 0$ satisfies*

$$-x \leq \frac{K}{(1 - \delta_A) a(\lambda)}. \tag{6}$$

Proof. Fix a task draw for which Lagrangian optimality holds, and let z be the selected continuation. Put $A := A(\lambda)$ and $Z := A^c$. After promise keeping has been dualized, Lagrangian optimality implies that the selected matching and continuation weakly dominate every alternative satisfying matching feasibility and (PC_{ik}).

Fix an assigned value $x < 0$ and let i be its recipient.

First suppose $\lambda_i > 0$. Remove this one negative assignment and retain the continuation promise. The modified decision satisfies matching feasibility and (PC_{ik}). Optimality gives $(1-\delta_P)+(1-\delta_A)\lambda_i x \geq 0$. Therefore

$$-x \leq \frac{1-\delta_P}{(1-\delta_A)a(\lambda)} \leq \frac{K}{(1-\delta_A)a(\lambda)}.$$

Now let $I^- \subseteq Z$ contain all zero-costate recipients of negative tasks, put $d := |I^-|$, and write $h_j := \ell(x_j)$ for $j \in I^-$. Remove every assignment to I^- and retain every other assignment. Apply part (i) of Lemma 13 to z_A, and let $b_A \in \mathcal{B}_A$ be the resulting completion. Because $|A| \geq K$, its total is $b_A(A) = r_K$. Set $b_Z = 0$; the resulting vector b belongs to $\mathcal{U}$, since for every $S \subseteq [N]$, $b(S) = b_A(S \cap A) \leq r_{|S \cap A|} \leq r_{|S|}$. Replace z by b. Constraint (PC_{ik}) holds for active workers because $b_A \geq z_A$. Every retained assignment to a zero-costate worker is nonnegative and has floor zero, so $b_Z = 0$ satisfies (PC_{ik}). The modified matching and continuation satisfy matching feasibility and (PC_{ik}).

The removed tasks have zero current λ-weight. Optimality for this realized task vector gives

$$\delta_A \lambda_A \cdot (b_A - z_A) \leq d(1-\delta_P) + \delta_P\{\Phi(z) - \Phi(b)\} \leq K.$$

The last inequality uses $d \leq K$ and $0 \leq \Phi \leq K$. Moreover,

$$\lambda_A \cdot (b_A - z_A) \geq a(\lambda)\{r_K - z(A)\} \geq a(\lambda) z(Z) \geq a(\lambda) \sum_{j \in I^-} h_j.$$

The middle inequality uses $z([N]) \leq r_K$, and the last uses (PC_{ik}). Combining the displays and using $\delta_A h_j = (1-\delta_A)(-x_j)$ gives

$$K \geq (1-\delta_A)a(\lambda) \sum_{j \in I^-} (-x_j).$$

The claimed bound follows for every $j \in I^-$. $\square$

Proof of Theorem 3(i). Attach the costate path supplied by Lemma 12 to the arbitrary optimal contract fixed at the start of the subsection. By Proposition 11, the stopping time

$$T_K := \inf\{t : |A(\lambda_t)| \geq K\}$$

is finite almost surely. At a date $t \geq T_K$, put $a_t := \min_{i \in A(\lambda_t)} \lambda_{i,t}$. By Lemma 21, a negative assignment requires at least one of the K task draws to belong to $[-K/((1-\delta_A)a_t), 0)$. The density bound $\bar{f}$ and a union bound give

$$\Pr(\text{a negative task is assigned at } t \mid \mathcal{F}_t) \leq \min\left\{1, \frac{K^2 \bar{f}}{(1-\delta_A)a_t}\right\}. \tag{7}$$

Partition the dates $t \geq T_K$ into intervals on which the set of positive costates is constant. There are at most $N - K + 1$ such intervals. On any one of them, denoted $[\tau, \sigma)$, (3) gives $a_t \geq q^{t-\tau} a_\tau$ for $\tau \leq t < \sigma$. Taking conditional expectations, applying (7), and summing gives

$$\mathbb{E}_\tau \left[\sum_{t=\tau}^{\sigma-1} \mathbf{1}\{\text{a negative task is assigned at } t\} \right] \leq \frac{K^2 \bar{f}}{(1-\delta_A) a_\tau} \sum_{n=0}^{\infty} q^{-n} < \infty.$$

Thus each interval contains only finitely many dates with a negative assignment, almost surely. There are finitely many intervals after T_K and finitely many dates before it, so negative-valued tasks are assigned at only finitely many dates almost surely.

Every positive task is assigned, and atomlessness excludes zero-valued tasks almost surely. Thus, from some finite date onward, the contract assigns exactly the positive-valued tasks. The attached costates did not change the contract's choices, so this conclusion applies to the arbitrary optimal contract fixed at the start of the subsection. □

A.5.2 The case $q < 1$

This subsection proves part (iii) of Theorem 3 in three steps. For any costate process constructed in Lemma 12, write $L_t := \sum_{i=1}^{N} \lambda_{i,t}$ for the total costate. First, we show that L_t is small with positive lower frequency. A uniform bound on expected multipliers prevents L_t from remaining large, while runs of positive tasks set all multipliers to zero and reduce L_t by the factor $q < 1$ each period. Consequently, dates with small L_t occur together with any positive-probability set of task vectors at positive lower frequency.

Second, we construct a positive-probability set B_K of all-negative task vectors such that, whenever L_t is sufficiently small, every optimal one-period decision assigns all K tasks and the participation multipliers have a uniformly positive total. The observation above that every positive task is assigned gives the positive-task conclusion. The first two steps show that negative tasks are assigned with positive lower frequency under every optimal Markov contract.

The remaining conclusions concern recipients and therefore require a particular optimal contract. We select an optimal assignment in which a worker with a higher costate receives weakly better current utility, with idleness corresponding to zero utility. On the recurring dates for which $V_t \in B_K$ and L_t is sufficiently small, the worst task goes to a worker with a smallest current costate. Symmetry and concavity order next-period costates with next-period promises. Combining this order with the strictly ordered participation floors and the costate equation shows that the worst-task recipient has a multiplier weakly larger than every other worker's and a weakly largest next-period costate. The positive lower bound on the total multiplier makes this worker's multiplier strictly positive. Whenever any costate remains zero, the worst-task recipient has zero costate, so each such date makes another costate positive. Because these dates recur, every worker's costate becomes positive in finite time.

Why the total costate is frequently small. Fix an optimal Markov contract and attach the costates and minimum-norm multipliers supplied by Lemma 12. Write ℓ_t for the selected matching's vector of participation floors. The process satisfies (3); in particular, $\mu_{i,t} = 0$ whenever $\ell_{i,t} = 0$.

The first lemma bounds the expected increase in costates due to binding participation constraints. Its factor $1 - \delta_P$ is the principal's current payoff from one additional assignment.

Lemma 22 (Uniform bound on expected multipliers)**.** *There is a constant $C_\mu < \infty$, depending only on the primitives, such that*

$$\mathbb{E}\left[\sum_{i=1}^{N} \mu_{i,t} \mid \mathcal{F}_t\right] \leq C_\mu(1 - \delta_P)$$

almost surely for every finite t.

Proof. Let $H := \ell(V)$. On $(0, \bar{h})$, this random variable has density

$$g_H(h) = \frac{\delta_A}{1 - \delta_A} f\left(-\frac{\delta_A}{1 - \delta_A} h\right).$$

Write $\bar{g} := \delta_A \bar{f}/(1 - \delta_A)$ for an upper bound on this density. The atom at $H = 0$ contributes no multiplier because Lemma 12 sets the multiplier on a zero floor equal to zero.

Condition on $\mathcal{F}_t$. Fix a matching ϱ, a pair (i, k) with $\varrho_{ik} = 1$, and all task values other than V_k. Let ℓ^{-ik} be the floor vector after deleting this assignment; its i-th coordinate is zero. If ϱ is feasible for some negative value of V_k, then $I_{\varrho ik} := \{h \geq 0 : \ell^{-ik} + he_i \in \mathcal{U}\}$ is an interval containing zero. Otherwise, this assignment contributes nothing to the bound. For $h \in I_{\varrho ik}$, define

$$D_{\varrho ik}(h) := \Gamma_{\lambda_t}(\ell^{-ik}) - \Gamma_{\lambda_t}(\ell^{-ik} + he_i).$$

Concavity and monotonicity of Γ_{λ_t} make $D_{\varrho ik}$ finite, nonnegative, nondecreasing, and convex, with $D_{\varrho ik}(0) = 0$. The supporting-slope inequality used to select μ_t, applied while varying only h, gives $D_{\varrho ik}(h') \geq D_{\varrho ik}(h) + \mu_{i,t}(h' - h)$. Hence $\mu_{i,t} = D'_{\varrho ik}(h)$ for almost every selected $h > 0$. The almost-everywhere gradient identity in Lemma 12, together with Fubini's theorem, ensures that, for almost every configuration of the other task values, this identity holds for Lebesgue-almost every h at which ϱ is selected.

If ϱ is selected at $h > 0$, then $h = \ell(V_k) = -(1-\delta_A)V_k/\delta_A$. Comparing ϱ with the matching that deletes the assignment of task k in (2) gives

$$0 \le (1-\delta_P) - \delta_A \lambda_{i,t} h - D_{\varrho ik}(h).$$

Since $\lambda_{i,t} \ge 0$, every selected h satisfies $D_{\varrho ik}(h) \le 1-\delta_P$. Let $S_{\varrho ik} := \{h \in I_{\varrho ik} \cap (0, \bar{h}) : \varrho \text{ is selected}\}$. Since the derivative is nonnegative,

$$\int_{S_{\varrho ik}} D'_{\varrho ik}(h)\,dh \le \int_{\{h \in I_{\varrho ik} : D_{\varrho ik}(h) \le 1-\delta_P\}} D'_{\varrho ik}(h)\,dh \le 1-\delta_P.$$

The sublevel set in the second integral is an interval beginning at zero. Absolute continuity gives the last inequality when its upper endpoint is in the interior of $I_{\varrho ik}$. At a boundary endpoint, apply the same argument on increasing compact subintervals and use monotone convergence.

Multiplying by $g_H(h) \le \bar{g}$ and integrating over the other task values shows that each fixed (ϱ, i, k) contributes at most $\bar{g}(1-\delta_P)$ to the conditional expectation. Moreover,

$$\sum_i \mu_{i,t} = \sum_{\varrho \in \mathcal{M}^{N,K}} \sum_{(i,k):\varrho_{ik}=1} \mu_{i,t} \mathbf{1}\{m_t = \varrho,\ V_{k,t} < 0\} \quad \text{almost surely,}$$

because every worker without an assigned negative task has a zero floor and hence a zero multiplier. A matching assigns at most K tasks, so the result holds with $C_\mu = K|\mathcal{M}^{N,K}|\bar{g}$. □

The multiplier bound implies that, above a fixed threshold, expected L_{t+1}, conditional on the current history, is at most a fixed fraction of L_t. After L_t falls below that threshold, a run of positive tasks makes it arbitrarily small because all multipliers vanish during the run. The next lemma shows that dates with a small total costate and a task vector in any given positive-probability set occur with positive lower frequency.

Lemma 23 (Dates with small total costate have positive frequency)**.** *Suppose $q < 1$. If $\eta > 0$ and $B \subseteq [\underline{v}, \overline{v}]^K$ is Borel with $F^K(B) > 0$, then*

$$\liminf_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{1}\{L_t \le \eta,\ V_t \in B\} > 0 \qquad \textit{almost surely.}$$

Proof. Taking sums and conditional expectations in (3) and applying Lemma 22 gives

$$\mathbb{E}[L_{t+1} \mid \mathcal{F}_t] \le qL_t + C_0, \qquad C_0 := \frac{C_\mu(1-\delta_P)}{\delta_P}.$$

Choose $M > C_0/(1-q)$ and set $\rho := q + C_0/M < 1$. Whenever $L_t > M$, the preceding bound gives $\mathbb{E}[L_{t+1} \mid \mathcal{F}_t] \le \rho L_t$. For a finite $\{\mathcal{F}_t\}$-stopping time σ, let $\tau_\sigma := \inf\{t \ge \sigma : L_t \le M\}$. Induction, applying the contraction only before τ_σ, gives

$$\mathbb{E}[L_{\sigma+n}\mathbf{1}\{\tau_\sigma > \sigma + n\} \,|\, \mathcal{F}_\sigma] \le \rho^n L_\sigma.$$

Since $L_{\sigma+n} > M$ on the event inside the indicator,

$$\mathbb{P}(\tau_\sigma > \sigma + n \mid \mathcal{F}_\sigma) \le \frac{L_\sigma}{M}\rho^n. \tag{8}$$

Thus $\tau_\sigma < \infty$ almost surely for every such σ.

Choose a nondegenerate closed interval $I_+ \subset (0, \overline{v})$, and let G_t be the event that every date-t task value lies in I_+. Write $p_G := \mathbb{P}(G_t) > 0$. Every task is positive on G_t, so every task is assigned, every participation floor is zero, and the normalization in Lemma 12 gives $\mu_t = 0$ and $L_{t+1} = qL_t$.

Choose an integer $r \ge 0$ with $q^r M \le \eta$, and put $d := r + 1$. Define

$$\sigma_1 := \inf\{t \ge 0 : L_t \le M\}, \qquad \sigma_{n+1} := \inf\{t \ge \sigma_n + d : L_t \le M\}.$$

Since L_t is $\mathcal{F}_t$-measurable, these are stopping times; (8) makes them finite almost surely. Iterating the conditional expectation bound over the next d dates gives

$$\mathbb{E}[L_{\sigma_n+d} \mid \mathcal{F}_{\sigma_n}] \le D := q^d M + C_0 \frac{1-q^d}{1-q}.$$

Applying (8) from date $\sigma_n + d$ and then conditioning on $\mathcal{F}_{\sigma_n}$ yields

$$\mathbb{P}(\sigma_{n+1} - \sigma_n > d + m \mid \mathcal{F}_{\sigma_n}) \le \frac{D}{M}\rho^m \qquad (m \ge 0).$$

Set $\mathcal{H}_n := \mathcal{F}_{\sigma_n}$ and $\Delta_n := \sigma_{n+1} - \sigma_n$. The geometric tail bound gives uniformly bounded conditional second moments and a constant $\overline{d} < \infty$ such that $\mathbb{E}[\Delta_n \mid \mathcal{H}_n] \le \overline{d}$. Since Δ_n is $\mathcal{H}_{n+1}$-measurable, the centered gaps $\Delta_n - \mathbb{E}[\Delta_n \mid \mathcal{H}_n]$ are martingale differences. Their strong law gives

$$\frac{1}{n}\sum_{j=1}^{n-1} \big(\Delta_j - \mathbb{E}[\Delta_j \mid \mathcal{H}_j]\big) \longrightarrow 0 \qquad \text{almost surely.}$$

Since $\sigma_n = \sigma_1 + \sum_{j=1}^{n-1} \Delta_j$, it follows that

$$\limsup_{n\to\infty} \frac{\sigma_n}{n} \leq \overline{d} \qquad \text{almost surely.} \tag{9}$$

Let A_n be the event that G_t occurs at $t = \sigma_n, \ldots, \sigma_n + r - 1$ and that $V_{\sigma_n + r} \in B$. Conditional on $\mathcal{H}_n$, these $r + 1$ task vectors are independent and each has law F^K, because σ_n is an $\{\mathcal{F}_t\}$-stopping time. Hence

$$\mathbb{P}(A_n \mid \mathcal{H}_n) = p_G^r F^K(B) =: p_* > 0.$$

Because $d = r + 1$, $A_n \in \mathcal{F}_{\sigma_n + d} \subseteq \mathcal{H}_{n+1}$. Thus $\mathbf{1}_{A_n} - p_*$ is a bounded martingale difference, and

$$\frac{1}{n}\sum_{j=1}^{n} \mathbf{1}_{A_j} \longrightarrow p_* \qquad \text{almost surely.}$$

On A_n, the r positive-task dates give $L_{\sigma_n + r} = q^r L_{\sigma_n} \leq \eta$, and $V_{\sigma_n + r} \in B$. Put $t_n := \sigma_n + r$ and $N(T) := \max\{n : t_n < T\}$, with $N(T) = 0$ if the set is empty. On the probability-one event on which the preceding two limits hold, $N(T) \to \infty$, and $T \leq t_{N(T)+1}$ implies

$$\liminf_{T\to\infty} \frac{N(T)}{T} \geq \frac{1}{\overline{d}}.$$

Each successful A_j with $j \leq N(T)$ contributes at the distinct date $t_j < T$. Therefore, for all sufficiently large T,

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbf{1}\{L_t \leq \eta,\ V_t \in B\} \geq \frac{N(T)}{T} \frac{1}{N(T)} \sum_{j=1}^{N(T)} \mathbf{1}_{A_j}.$$

Taking the lower limit gives $p_*/\overline{d} > 0$. □

A recurring set of all-negative tasks. We apply the preceding result to a set B_K on which every task is negative. The next lemma constructs this set so that, at a sufficiently small total costate, assigning all K tasks strictly dominates every incomplete matching. It also shows that the resulting participation multipliers have a uniformly positive total. Let ρ_K be the largest common promise on the plateau, as defined in Lemma 10, and set $c_K := \overline{V}_K/K = \mathbb{E}[[V]_+]$. Under the maintained condition $\mathbb{E}[V] < \bar{h}$, that lemma gives $\rho_K < \min\{\bar{h}, c_K\}$.

Lemma 24 (Full assignment and positive multipliers when total costate is small)**.** *Suppose $\mathbb{E}[V] < \bar{h}$. There are constants $\eta_0 > 0$ and $\kappa > 0$, and a compact set $B_K \subset (\underline{v}, 0)^K$ with $F^K(B_K) > 0$, such that, whenever $L_t \leq \eta_0$, every λ_t-optimal one-period decision assigns all K tasks at every $v \in B_K$. At each such v, the* *corresponding multiplier vector μ, selected as in Lemma 12, satisfies $\sum_i \mu_i \geq \kappa$.*

Proof. Let ω be a modulus of continuity for $\Phi^{N,K}$ on its compact domain, using the sup norm. By the separation just established, choose $\epsilon > 0$ such that

$$\rho_K + \epsilon < \min\{\bar{h}, c_K\}, \qquad \delta_P \omega(\epsilon) < 1 - \delta_P.$$

Choose nondegenerate, pairwise disjoint closed intervals $J_1 < \cdots < J_K$ whose union lies in $(\rho_K, \rho_K + \epsilon)$, where every point of J_r is smaller than every point of J_{r+1}. Let B_K be the set of task vectors for which the r-th smallest participation floor lies in J_r for every r. The ordered-floor map is continuous, so B_K is compact. Every task in B_K is negative, and strict positivity of the density gives $F^K(B_K) > 0$. Compactness and strict separation of the intervals give positive lower bounds on the floors' distance from ρ_K and from one another.

Fix a task vector in B_K and any set S of K recipients. The vector ℓ that gives these recipients their task floors is feasible. Indeed, a set of s workers contains at most $\min\{s, K\}$ positive coordinates, each below c_K. For $s \leq K$, the sum of the s largest positive task values is pointwise at least s/K times the sum of all K positive task values, so $sc_K \leq \overline{V}_s$. For $s > K$, $Kc_K = \overline{V}_K = \overline{V}_s$. Every subset constraint therefore has strict slack, and the positive-floor vector lies in the interior of its feasible set. Moreover, $\|\ell - \rho_K \mathbf{1}_S\|_\infty < \epsilon$, and $\rho_K \mathbf{1}_S \in \mathcal{P}$.

At zero costate, assigning all K tasks and continuing at ℓ gives an objective of at least $K(1 - \delta_P) + \delta_P[\Phi(0) - \omega(\epsilon)]$. Any incomplete matching has objective at most $(K-1)(1-\delta_P) + \delta_P \Phi(0)$. The full matching therefore has an advantage of at least

$$\gamma := (1 - \delta_P) - \delta_P \omega(\epsilon) > 0.$$

To extend this strict comparison to small costates, put

$$C_\lambda := (1 - \delta_A) \max\{|\underline{v}|, \overline{v}\} + \delta_A \overline{V}_1.$$

For every feasible matching, the absolute current costate term is at most $(1 - \delta_A) \max\{|\underline{v}|, \overline{v}\} L_t$, while

$$0 \leq \Gamma_{\lambda_t}(\ell) - \Gamma_0(\ell) \leq \delta_A \overline{V}_1 L_t.$$

Consequently, $|W_{\lambda_t}(m; v) - W_0(m; v)| \leq C_\lambda L_t$ for every feasible matching. Choose $\bar{\eta} > 0$ so that $2C_\lambda \bar{\eta} < \gamma$. Every λ_t-optimal matching then assigns all K tasks whenever $L_t \leq \bar{\eta}$ and $v \in B_K$.

After reducing $\bar{\eta}$ if necessary, the total multiplier is uniformly bounded away from zero. Otherwise, for every integer n, one could choose a state u^n, a costate $\lambda^n \in \Lambda(u^n)$ with $\sum_i \lambda_i^n \leq 1/n$, a task vector in B_K, a recipient set S^n, a continuation z^n, and minimum-norm multipliers μ^n selected as in Lemma 12 such that $\sum_i \mu_i^n \leq 1/n$. Pass to a subsequence with a fixed recipient set S. Compactness gives $\ell^n \to \ell^*$ and $z^n \to z^*$. Equation (3) gives $\lambda^{n,+} \to 0$. Since $\lambda^{n,+} \in \Lambda^{N,K}(z^n)$, closedness of the subgradient graph implies $0 \in \Lambda^{N,K}(z^*)$, and hence $z^* \in \mathcal{P}$.

Constraint (PC_{ik}) gives $z^* \geq \ell^*$, and every coordinate of ℓ^* on S lies a fixed positive distance above ρ_K. Let $h := \min_{i \in S} \ell_i^* > \rho_K$. Downward closure gives $h\mathbf{1}_S \in \mathcal{P}$, and symmetry gives $h\mathbf{1}_T \in \mathcal{P}$ for every set T of K workers. This contradicts the definition of ρ_K. The contradiction supplies $\kappa > 0$; take $\eta_0 \leq \bar{\eta}$ small enough for this bound to hold. □

Selecting assignments ordered by costates. The preceding lemmas give the negative-task conclusion for every optimal Markov contract. The conclusions about recipients use a selected optimal contract whose matching is ordered by current costates. We call workers with the same costate an *equal-costate group*. An idle worker receives zero current utility, so the same ordering compares assigned tasks with idleness.

Lemma 25 (Assignments ordered by current costates)**.** *There exists an optimal recursive contract with a measurable finite costate process and the minimum-norm multipliers from Lemma 12. After almost every task draw at every reached history, $\lambda_{i,t} < \lambda_{j,t}$ implies $Y_{i,t} \leq Y_{j,t}$, where $Y_{i,t} := Y_i(m_t, V_t)$ and $Y_{i,t} = 0$ when worker i is idle. Thus, when all K tasks are negative and assigned, each recipient's costate is weakly below each idle worker's costate, and the worst task goes to a worker with a minimal current costate.*

Proof. Represent each idle worker by a placeholder with task value and participation floor zero. Whenever a worker with a larger costate receives a lower-valued task than a worker with a smaller costate, apply the exchange from Lemma 6. Each exchange preserves both workers' realized utility delivery and the principal's Bellman payoff. Repeating the exchange produces the stated ordering after finitely many steps. Choosing at each step the first offending pair according to a fixed ordering makes the procedure measurable. In particular, each recipient of a negative task has a costate weakly below each idle worker's costate, and the worst task goes to a worker with a minimal costate.

At each pair (u, λ) with $\lambda \in \Lambda(u)$, use Lemmas 2 and 3 to select measurably a Bellman-optimal one-period rule feasible at u. Lagrangian optimality makes this rule λ-optimal after almost every task vector. Apply the finite measurable exchange procedure above, and then use Lemma 12 to select the minimum-norm multiplier and next costate. The exchanges preserve the delivered-utility vector, so the modified rule remains feasible and optimal at u. The local construction in Lemma 12 makes these last two choices measurable. Iterating the combined selection from $(u_0, \lambda_0) = (0, 0)$ gives the required optimal recursive contract and costate process. $\square$

Proof of Theorem 3(iii). Fix any optimal Markov contract and attach the costates and multipliers from Lemma 12. Work on the probability-one event on which this construction and Lagrangian optimality hold at every finite date. The construction leaves the contract decisions unchanged.

Positive and negative assignments. Every optimal contract assigns every positive task after almost every draw. The events on which all K tasks are positive are i.i.d. and have strictly positive probability. The strong law of large numbers therefore implies that they have positive long-run frequency, and every task is assigned on each of these dates.

Lemma 23, applied with η_0 and B_K, shows that dates satisfying $L_t \leq \eta_0$ and $V_t \in B_K$ have positive lower frequency. All K tasks are negative on these dates, and Lemma 24 shows that they are assigned. Because the attached construction leaves the contract unchanged, these task-level conclusions hold for every optimal Markov contract.

For the conclusions about recipients, use the optimal contract and costate process selected in Lemma 25, starting from $(u_0, \lambda_0) = (0, 0)$.

Movement of the worst-task recipient. When $V_t \in B_K$, all K tasks are assigned if $L_t \leq \eta_0$. List their recipients as $a_1, \ldots, a_K$, ordered first by increasing costate and, within a costate tie, by decreasing participation floor. The ordering in Lemma 25 gives

$$\lambda_{a_1,t} \leq \cdots \leq \lambda_{a_K,t}, \qquad \ell_{a_1,t} > \cdots > \ell_{a_K,t} > 0,$$

where a_1 receives the worst task and belongs to the current equal-costate group with the lowest costate. The strict floor inequalities come from the disjoint intervals used to define B_K.

The corresponding multipliers satisfy $\mu_{a_1,t} \geq \cdots \geq \mu_{a_K,t}$. To see this, suppose that $\mu_{a_r,t} < \mu_{a_s,t}$ for some $r < s$. Then $\mu_{a_s,t} > 0$, so complementary slackness gives $u_{a_s,t+1} = \ell_{a_s,t}$, whereas $u_{a_r,t+1} \geq \ell_{a_r,t} > \ell_{a_s,t}$. By Proposition 6, a supporting costate is ordered with the promise vector. Hence $\lambda_{a_r,t+1} \geq \lambda_{a_s,t+1}$. The costate recursion instead gives

$$\lambda_{a_r,t+1} - \lambda_{a_s,t+1} = q(\lambda_{a_r,t} - \lambda_{a_s,t}) + \frac{\mu_{a_r,t} - \mu_{a_s,t}}{\delta_P} < 0,$$

a contradiction. Since all nonrecipients have zero floor and hence zero multiplier, Lemma 24 now implies

$\mu_{a_1,t} \geq \kappa/K$.

It follows that $\lambda_{a_1,t+1} \geq \kappa/(K\delta_P)$. Every worker with zero multiplier has next costate at most qL_t. If recipient a_r, with $r > 1$, has a positive multiplier, complementary slackness and the strict floor order give $u_{a_1,t+1} > u_{a_r,t+1}$; the same ordering property gives $\lambda_{a_1,t+1} \geq \lambda_{a_r,t+1}$. Choose

$$\eta_1 < \min\left\{\eta_0, \frac{\kappa}{qK\delta_P}\right\}.$$

At a date with $V_t \in B_K$ and $L_t \leq \eta_1$, worker a_1 therefore has a maximal next costate and strictly exceeds every worker with zero multiplier. Thus a_1 belongs to the next period's equal-costate group with the highest costate.

Lemma 23, applied with η_1 and B_K, shows that dates satisfying $L_t \leq \eta_1$ and $V_t \in B_K$ have positive lower frequency almost surely.

Eventually positive costates. Let $Z_t := \{i : \lambda_{i,t} = 0\}$. Once a costate becomes positive, (3) keeps it positive. If $Z_t \neq \emptyset$ at one of these dates, the group with the lowest costate is exactly Z_t, so $a_1 \in Z_t$. The argument above gives this worker a multiplier of at least κ/K, so its next costate is positive. Before all costates are positive, each such date reduces $|Z_t|$ by at least one. These dates occur infinitely often, so after at most N such dates all N costates are positive; the last of those dates is finite almost surely. □

A.5.3 Equal patience: bounded costates and limiting assignments

Set $\delta_A = \delta_P = \delta$. Then (3) becomes $\lambda_{t+1} = \lambda_t + \mu_t/\delta$, so every coordinate of the selected costate is weakly increasing. This subsection proves part (ii) of Theorem 3 in three steps. First, we show that the costate vector nevertheless remains bounded and therefore converges to a finite, generally history-dependent vector λ_∞. Second, along almost every realized path, boundedness gives a fixed interval just below zero such that all tasks are assigned whenever every task value lies in that interval. The strong law gives positive frequency both to these all-negative draws and to draws in which all task values lie in a fixed positive interval. Third, convergence of the costates makes the one-period assignment objective converge uniformly to the objective evaluated at λ_∞. Since there are only finitely many matchings, for F^K-almost every task vector v, the date-t rule eventually selects a matching in $\mathcal{A}_{\lambda_\infty}(v)$. Integrating over an independent task draw gives the convergence stated in the theorem.

Why costates remain bounded. Every strict increase of the running maximum to a value above a fixed threshold must bring at least one coordinate across that threshold from below. Each coordinate can cross only once, which bounds the running maximum and hence the costate vector.

Lemma 26 (Costates are bounded under equal patience)**.** *For the costate process constructed in Lemma 12, if $q = 1$, then $\sup_t \|\lambda_t\|_\infty < \infty$ almost surely. Consequently, λ_t converges coordinatewise to a finite vector λ_∞.*

Proof. By Proposition 6, the plateau contains a nonzero vector. Convexity and symmetry imply that the average of its coordinate permutations also belongs to the plateau. Hence $b\mathbf{1} \in \mathcal{P}$ for some $b > 0$. Put $D := (1-\delta)/\delta$, choose $L > D/b$, and let $M_t := \max_i \lambda_{i,t}$.

Each coordinate of λ_t, and hence M_t, is nondecreasing. Consider a date at which $M_{t+1} > \max\{M_t, L\}$, and let $B := \{i : \lambda_{i,t+1} = M_{t+1}\}$. For every $i \in B$, $\lambda_{i,t} \leq M_t < M_{t+1}$, so the costate recursion gives $\mu_{i,t} > 0$. The zero-floor normalization then implies that worker i receives a negative task x_i, and complementarity gives $u_{i,t+1} = \ell(x_i) = D(-x_i) > 0$.

Suppose $\lambda_{i,t} \geq L$ for every $i \in B$. Fix $i \in B$, delete only that worker's assignment, and retain the selected continuation. The modified matching and continuation remain feasible for this task vector. Their continuation term is unchanged, so λ_t-optimality after this task vector gives $0 \leq (1-\delta)(1+\lambda_{i,t}x_i)$. Therefore $-x_i \leq 1/\lambda_{i,t}$ and $u_{i,t+1} \leq D/L < b$. Since i was arbitrary, this bound holds for every $i \in B$. Because $-\Phi$ is symmetric and convex, every supporting costate satisfies $(u_{i,t+1} - u_{j,t+1})(\lambda_{i,t+1} - \lambda_{j,t+1}) \geq 0$. For $i \in B$ and $j \notin B$, the second factor is positive, so $u_{j,t+1} \leq u_{i,t+1} < b$. Thus $u_{t+1} < b\mathbf{1}$ coordinatewise.

Downward closure now puts u_{t+1} in $\mathcal{P}$. For every coordinate j, choose $\varepsilon_j > 0$ with $u_{t+1} + \varepsilon_j e_j \leq b\mathbf{1}$. This raised vector also lies in $\mathcal{P}$. Since both vectors lie on the plateau, the subgradient inequality for $-\Phi$ gives $0 \geq \lambda_{j,t+1}\varepsilon_j$. Hence every next costate is zero, contradicting $M_{t+1} > L$.

Every strict increase of M_t above L therefore includes a coordinate that crosses L from below. Each coordinate can cross L at most once, so there are at most N such increases. Values at or below L are bounded, and every value produced by one of these finitely many increases is finite. Hence M_t is bounded. Coordinatewise monotonicity then gives the finite limit λ_∞. $\square$

Why tasks just below zero are assigned. A task just below zero requires only a small continuation promise. The next lemma makes this observation uniform over every costate vector in a bounded set.

Lemma 27 (Assignment near zero under bounded costates)**.** *Suppose $K \leq N$ and $\delta_A = \delta_P = \delta$. For every finite $M \geq 0$, there is $\varepsilon_M \in (0, -\underline{v})$ such that, for every $\lambda \in \mathbb{R}^N_+$ with $\|\lambda\|_1 \leq M$ and every $v \in (-\varepsilon_M, 0)^K$, every λ-optimal matching assigns all K tasks.*

Proof. Put $D := (1-\delta)/\delta$. Choose $0 < a < \min_{1\leq r\leq N} \overline{V}_r/r$, so $a\mathbf{1} \in \mathcal{U}$, and write $A_\lambda(z) := \delta\Phi(z) + \delta\lambda \cdot z$.

Fix $\lambda \in \mathbb{R}^N_+$ with $\|\lambda\|_1 \leq M$, and let z^0 maximize A_λ on $\mathcal{U}$. If $0 \leq \ell \leq h\mathbf{1}$ with $0 \leq h \leq a$, put $\alpha := h/a$ and $\widehat{z} := (1-\alpha)z^0 + \alpha a\mathbf{1}$. Convexity gives $\widehat{z} \in \mathcal{U}$, and $\widehat{z} \geq \alpha a\mathbf{1} = h\mathbf{1} \geq \ell$. Hence $\Gamma_\lambda(\ell) \geq A_\lambda(\widehat{z})$, and concavity gives

$$0 \le \Gamma_\lambda(0) - \Gamma_\lambda(\ell) \le \frac{h}{a}\{A_\lambda(z^0) - A_\lambda(a\mathbf{1})\}.$$

Every coordinate of z^0 is at most $\overline{V}_1$, while $0 \le \Phi \le K$ and $A_\lambda(a\mathbf{1}) \ge 0$. Therefore

$$\Gamma_\lambda(0) - \Gamma_\lambda(\ell) \le C_M h, \qquad C_M := \frac{\delta(K + M\overline{V}_1)}{a}.$$

Fix $\varepsilon > 0$ with $D\varepsilon \le a$, take $v \in (-\varepsilon, 0)^K$, and assign the K tasks to distinct workers. If ℓ is the resulting floor vector, then $\lambda \cdot Y \ge -M\varepsilon$ and $0 \le \ell \le D\varepsilon\mathbf{1}$. Downward closure gives $\ell \in \mathcal{U}$, so this matching has objective at least

$$(1-\delta)K - (1-\delta)M\varepsilon + \Gamma_\lambda(0) - C_M D\varepsilon.$$

Any matching with at most $K-1$ tasks has objective at most $(1-\delta)(K-1) + \Gamma_\lambda(0)$, because all task values are negative and $\lambda \ge 0$, while its floor restriction can only reduce Γ_λ. Choose

$$0 < \varepsilon_M < \min\left\{-\underline{v}, \frac{a}{D}, \frac{1-\delta}{(1-\delta)M + C_M D}\right\}.$$

This matching strictly dominates every matching that assigns fewer tasks. Hence every λ-optimal matching assigns all K tasks. □

Proof of Theorem 3(ii). Fix an arbitrary optimal Markov contract starting from $u_0 = 0$, and attach the costates and multipliers from Lemma 12; this leaves every contract decision unchanged. On a common probability-one event, the costate equation and Lagrangian optimality hold at every realized finite date. Work on this event. The costate conclusion follows from Lemma 26.

Positive and negative assignments recur. Fix a nondegenerate compact interval $I_+ \subset (0, \overline{v})$. Every positive task is assigned. The events $\{V_t \in I_+^K\}$ are i.i.d. with positive probability, so the strong law gives them positive long-run frequency.

For each positive integer M, let $E_M^- := (-\varepsilon_M, 0)^K$, with ε_M from Lemma 27, and put $H_M := \{\sup_t \|\lambda_t\|_1 \le M\}$. On H_M, every occurrence of E_M^- is fully assigned. Intersecting the common event above with the strong-law events for this countable family shows that, for every M, the dates on which $V_t \in E_M^-$ have limiting frequency $\mathbb{P}(V \in (-\varepsilon_M, 0))^K > 0$. Since N is finite, Lemma 26 places each path in some H_M. This proves the two frequency conclusions.

The rule approaches the limiting set of optimal matchings. It remains to compare the date-t assignment problem with the problem evaluated at the limiting costate. Compactness of $\mathcal{U}$, bounded task values, and (2) give a constant $C < \infty$ such that, for all $\lambda, \widehat{\lambda} \in \mathbb{R}^N_+$,

$$\sup_{\substack{v \in [\underline{v}, \overline{v}]^K, \\ m \in \mathcal{M}^{N,K}(v)}} |W_\lambda(m; v) - W_{\widehat{\lambda}}(m; v)| \le C \|\lambda - \widehat{\lambda}\|_1.$$

Indeed, boundedness of $Y(m, v)$ controls the current linear term, while

$$|\Gamma_\lambda(\ell) - \Gamma_{\widehat{\lambda}}(\ell)| \le \delta \max_{z \in \mathcal{U}} |(\lambda - \widehat{\lambda}) \cdot z|$$

uniformly over feasible floors.

Fix a realized history on which $\lambda_t \to \lambda_\infty$. For each t, Lagrangian optimality gives an F^K-null set E_t outside which $m_t(v)$ maximizes $W_{\lambda_t}(\cdot; v)$. Let $E := \bigcup_t E_t$, which is also F^K-null. For every $v \notin E$, adding and subtracting W_{λ_t} gives

$$0 \le \max_{m \in \mathcal{M}^{N,K}(v)} W_{\lambda_\infty}(m; v) - W_{\lambda_\infty}(m_t(v); v) \le 2 \sup_{m \in \mathcal{M}^{N,K}(v)} |W_{\lambda_t}(m; v) - W_{\lambda_\infty}(m; v)|.$$

The right side converges to zero uniformly in v. Write $W^*_\infty(v) := \max_{m \in \mathcal{M}^{N,K}(v)} W_{\lambda_\infty}(m; v)$. For fixed $v \notin E$, let $\Delta(v)$ be the minimum of the strictly positive values in

$$\left\{W^*_\infty(v) - W_{\lambda_\infty}(m; v) : m \in \mathcal{M}^{N,K}(v)\right\}.$$

If this set contains no positive value, every feasible matching is a maximizer. Otherwise, finiteness of the matching set gives $\Delta(v) > 0$, and the preceding bound implies that $m_t(v) \in \mathcal{A}_{\lambda_\infty}(v)$ for all sufficiently large t. Thus

$$\mathbf{1}\{m_t(v) \notin \mathcal{A}_{\lambda_\infty}(v)\} \longrightarrow 0 \qquad \text{for every } v \notin E.$$

Dominated convergence now gives

$$F^K\{v : m_t(v) \notin \mathcal{A}_{\lambda_\infty}(v)\} \longrightarrow 0,$$

which is the independent-draw probability in the theorem.

The idle-worker deviation in the proof of Lemma 6 shows that every matching in $\mathcal{A}_{\lambda_\infty}(v)$ assigns every positive task. Taking an integer $M > \|\lambda_\infty\|_1$ in Lemma 27 also shows that every limiting maximizer fully assigns the fixed cube $(-\varepsilon_M, 0)^K$.

$\square$

Implication for the long-run assignment fraction. The three patience cases can now be combined to characterize the long-run fraction of available tasks that are assigned. Let M_t be the number of tasks assigned at date t, and let $Z_t := \sum_{k=1}^{K} \mathbf{1}\{V_{k,t} > 0\}$ be the number of positive tasks. Define $R_T := (TK)^{-1} \sum_{t=0}^{T-1} M_t$ and $p^+ := \Pr(V > 0)$. Every optimal Markov contract assigns every positive task, so $M_t \geq Z_t$ almost surely. The strong law gives

$$\frac{1}{TK} \sum_{t=0}^{T-1} Z_t \longrightarrow p^+ \qquad \text{almost surely.}$$

If $q > 1$, part (i) of Theorem 3 says that negative tasks are assigned at only finitely many dates. Atomlessness and countability exclude zero-valued tasks at every date on a probability-one event. Hence $M_t = Z_t$ from some finite date onward, so $R_T \to p^+$.

Suppose instead that $q \leq 1$, and let E_t^- be the event that some negative-valued task is assigned at date t. Since every positive task is assigned, $M_t \geq Z_t + \mathbf{1}\{E_t^-\}$. Parts (ii) and (iii) of Theorem 3 say that E_t^- has positive lower frequency. Therefore

$$\liminf_{T\to\infty} R_T \geq p^+ + \frac{1}{K} \liminf_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{1}\{E_t^-\} > p^+ \qquad \text{almost surely.}$$

A.5.4 Specialization to one agent

With one agent and one task, the promise set is $[0, V^+]$, the plateau is $[0, \bar{u}]$, and every positive costate supports a unique promise by Lemma 9. These properties translate the preceding costate results into the conclusions of Proposition 5.

Proof of Proposition 5. Specializing Lemma 12 to $N = K = 1$ attaches finite scalars $\lambda_t \in \Lambda(u_t)$ and $\mu_t \geq 0$ to the given optimal contract without changing its assignment or continuation rules. Write X_t for the realized assignment at date t. On a common probability-one event, these scalars satisfy at every finite date

$$\lambda_{t+1} = q\lambda_t + \frac{\mu_t}{\delta_P}, \qquad \mu_t\big(u_{t+1} - \ell(V_t)X_t\big) = 0,$$

where $\mu_t = 0$ when the realized participation floor is zero. Work on this event. By Proposition 4, the stopping time $\tau := \inf\{t : u_t > \bar{u}\}$ is finite almost surely. Every costate at a state above the plateau is positive. Hence $\lambda_\tau > 0$, and the recursion keeps $\lambda_t > 0$ for all $t \geq \tau$. Because the selected costate is finite at every finite date, Lemma 4 also implies that $u_t < V^+$ at every finite date.

Agent more patient: $q > 1$**.** Condition on a date-t pre-draw history with current promise $u := u_t \in (\bar{u}, V^+)$. The set $\Lambda(u)$ is a nonempty compact interval. Let $b := \max \Lambda(u) > 0$. For F-almost every task value, the contract's decision is b-optimal. Let r maximize $\delta_P \Phi(z) + \delta_A b z$ over $[0, V^+]$. Then $qb \in \Lambda(r)$, and r is unique because $qb > 0$. Monotonicity of the subgradients of $-\Phi$ gives $(qb - b)(r - u) \geq 0$, so $r \geq u$. Equality would put qb in $\Lambda(u)$, contradicting the definition of b. Hence $r > u$.

Let $a := \ell(V_t) X_t$ be the realized participation floor. If $a \leq r$, then r is feasible and uniqueness makes it the selected continuation. If $a > r$, feasibility requires $u_{t+1} \geq a > r$. Conditional on every such history, therefore, $u_{t+1} > u_t$ almost surely. Taking the intersection of these events over the countably many dates shows that the state increases strictly after τ.

The state is bounded above by V^+, so it converges. The recursion and $\mu_t \geq 0$ give $\lambda_t \geq q^{t-\tau} \lambda_\tau$, hence $\lambda_t \to \infty$. If the promise limit were below V^+, the tail of the path would lie in a compact subinterval of $(\bar{u}, V^+)$. Subgradients of the finite convex function $-\Phi$ are uniformly bounded on such an interval, contradicting $\lambda_t \in \Lambda(u_t)$ and $\lambda_t \to \infty$. Thus $u_t \to V^+$.

The scalar specialization of Lemma 21 says that every assigned negative value satisfies $-v \leq 1/((1-\delta_A)\lambda_t)$ for F-almost every v. Since F has positive density and the cutoff rule assigns almost every value above $\gamma(u_t)$, this implies $-\gamma(u_t) \leq 1/((1-\delta_A)\lambda_t)$. Otherwise an interval of assigned negative values with positive probability would violate the bound. Hence $\gamma(u_t) \to 0 = \gamma(V^+)$, proving part (i).

Equal patience: $q = 1$**.** The costate recursion makes λ_t nondecreasing, and Lemma 26 makes it bounded. Thus $\lambda_t \to \lambda_\infty < \infty$, with $\lambda_\infty \geq \lambda_\tau > 0$. At a state $u_t > \bar{u}$, the current positive costate supports only u_t. Since $q = 1$, the target continuation in Theorem 1 is therefore u_t, and $u_{t+1} = \max\{u_t, \ell(V_t) X_t\}$. The state is nondecreasing after τ, so it converges to some u_∞. Closedness of the subgradient graph gives $\lambda_\infty \in \Lambda(u_\infty)$. Moreover, $u_\infty > \bar{u}$ because $u_\tau > \bar{u}$, and $u_\infty < V^+$ by Lemma 4. Hence $u_\infty \in (\bar{u}, V^+)$.

It remains to identify the limiting cutoff. Write $X_t(v)$, $z_t(v)$, and $\mu_t(v)$ for the assignment, continuation, and multiplier maps at the date-t pre-draw history. Telescoping the equal-patience recursion along the realized path gives $\sum_{t \geq 0} \mu_t(V_t) = \delta(\lambda_\infty - \lambda_0) < \infty$. For each positive rational η, set $p_t(\eta) := F\{v : \mu_t(v) > \eta\}$. The events $\{\mu_t(V_t) > \eta\}$ occur only finitely often. Independence of the current draw gives $p_t(\eta) = \Pr(\mu_t(V_t) > \eta \mid \mathcal{F}_t)$, so the conditional Borel–Cantelli lemma gives $\sum_t p_t(\eta) < \infty$. In particular, $p_t(\eta) \to 0$. Take the countable intersection of the probability-one events on which this conclusion holds for every positive rational η.

Fix a realized history in this intersection. Taking the union over dates of the F-null exceptional sets, we may use $\lambda_t + \mu_t(v)/\delta \in \Lambda(z_t(v))$ for every t and every v outside one F-null set. For each $\varepsilon > 0$, compactness, closedness of the subgradient graph, and $\Lambda^{-1}(\lambda_\infty) = \{u_\infty\}$ give an $\alpha > 0$ such that $\lambda' \in \Lambda(z)$ and $|\lambda' - \lambda_\infty| < \alpha$ imply $|z - u_\infty| < \varepsilon$. Indeed, failure of this implication would produce another promise supported by λ_∞. For large t, $|\lambda_t - \lambda_\infty| < \alpha/2$. Choose a positive rational $\eta < \delta\alpha/2$. Then

$$F\{v : |z_t(v) - u_\infty| \geq \varepsilon\} \leq F\{v : \mu_t(v) > \delta\alpha/2\} \leq p_t(\eta) \longrightarrow 0.$$

Thus $z_t(\cdot) \to u_\infty$ in F-measure.

The convergence argument for W_{λ_t} in the proof of Theorem 3(ii) applies to the assignment maps. In the scalar problem, the assignment gain is strictly increasing in the task value and its tie set has F-measure zero. The limiting λ_∞-optimal rule therefore has a unique cutoff c_∞, and

$$F\{v : X_t(v) \neq \mathbf{1}\{v \geq c_\infty\}\} \longrightarrow 0.$$

The date-t rule has cutoff $\gamma(u_t)$. Because F has positive density, the preceding convergence implies $\gamma(u_t) \to c_\infty$: cutoffs separated by a fixed amount produce different assignments on an interval with positive probability.

The assignment and continuation maps are uniformly bounded, so their convergence in measure also gives convergence in $L^1(F)$. Since $\lambda_t > 0$ after τ, complementary slackness makes promise keeping bind. Passing to the limit gives $U(c_\infty) = u_\infty$. Passing to the limit in the Bellman equality gives

$$\Phi(u_\infty) = (1 - \delta)\Pr(V \geq c_\infty) + \delta\Phi(u_\infty).$$

To verify (PC_{ik}), take a subsequence along which both maps converge F-almost everywhere and pass to $z_t(v) \geq \ell(v)X_t(v)$. Continuity and monotonicity of ℓ give $\ell(c_\infty) \leq u_\infty$. Thus the cutoff c_∞, together with the constant continuation u_∞, is feasible and attains $\Phi(u_\infty)$. It is optimal at u_∞, so Theorem 1 gives $c_\infty = \gamma(u_\infty)$.

The equalities just established show that $\gamma(u_\infty)$ is feasible in the definition of $c^S(u_\infty)$. If $c^S(u_\infty) < \gamma(u_\infty)$, choose a feasible cutoff $c < \gamma(u_\infty)$ and keep the continuation fixed at u_∞. This gives the Bellman deviation

$$(1 - \delta_P)\Pr(V \geq c) + \delta_P\Phi(u_\infty) > (1 - \delta_P)\Pr(V \geq \gamma(u_\infty)) + \delta_P\Phi(u_\infty) = \Phi(u_\infty),$$

where strictness uses the positive density of F. This contradicts optimality. Thus $\gamma(u_\infty) = c^S(u_\infty)$, proving part (ii).

Principal more patient: $q < 1$**.** Write $\gamma_0 := \gamma(0) = \gamma(\bar{u})$ and $u^\sharp := \min\{V^+, \ell(\gamma_0)\}$. The proof of Proposition 4 gives a positive-probability set of values that are assigned on the plateau and have participation floors above $\bar{u}$. These values lie above γ_0 almost surely. Since ℓ is strictly decreasing on negative task values, $\ell(\gamma_0) > \bar{u}$. Together with $V^+ > \bar{u}$, this gives $u^\sharp > \bar{u}$.

We first show that the promise process remains in $[0, u^\sharp]$, and remains in $[\bar{u}, u^\sharp]$ after τ. The cutoff comparison in the proof of Theorem 1 gives $\gamma(u) \geq \gamma_0$ at every state. Hence every realized participation floor is at most both $\ell(\gamma_0)$ and V^+, and therefore at most $u^\sharp$. On the plateau, Theorem 1 selects either a continuation on the plateau or the participation floor, so the continuation is at most $u^\sharp$.

At a state $u > \bar{u}$, let $b := \min \Lambda(u) > 0$. The unique unconstrained continuation for the b-Lagrangian is the point r with $qb \in \Lambda(r)$. Monotonicity gives $(qb - b)(r - u) \geq 0$, hence $r \leq u$. Equality would put qb in $\Lambda(u)$, contradicting the definition of b, because $qb < b$. Thus $r < u$. The selected continuation is the larger of r and the participation floor. Starting from $u_0 = 0$, induction now gives $u_t \leq u^\sharp$ at every date.

Since $\lambda_\tau > 0$ and $\lambda_{t+1} \geq q\lambda_t$, the selected costate remains positive after τ. Only the zero costate supports a promise below $\bar{u}$, so $u_t \geq \bar{u}$ for every $t \geq \tau$. A positive costate can support $\bar{u}$ when $-\Phi$ has a kink there, so this lower bound can hold with equality.

The scalar specialization of Lemma 23, with its task-value set equal to the full support, gives, for every $\eta > 0$, $\lambda_t \leq \eta$ infinitely often almost surely. Small costates force the state close to the plateau. Indeed, for every $\varepsilon \in (0, u^\sharp - \bar{u})$, strict decline of Φ above the plateau gives $s_\varepsilon := (\Phi(\bar{u}) - \Phi(\bar{u} + \varepsilon))/\varepsilon > 0$. Convexity of $-\Phi$ implies that $u \geq \bar{u} + \varepsilon$ and $\lambda \in \Lambda(u)$ together require $\lambda \geq s_\varepsilon$. Thus every right neighborhood of $\bar{u}$ is visited infinitely often after τ.

Fix $x \in (\bar{u}, u^\sharp)$ and a neighborhood O of x whose closure lies in that interval. Because ℓ is continuous and strictly decreasing on negative task values, there is a compact interval $J \subset (\gamma_0, 0)$ with $F(J) > 0$ and $\ell(J) \subset O$. At zero costate, assignment strictly improves the one-period Lagrangian for every value in J. Continuity of this gain and compactness of J imply that every value in J remains strictly assigned whenever the current costate is sufficiently small.

Choose $\varepsilon > 0$ so that $\bar{u} + \varepsilon < \inf \ell(J)$, and then choose $\eta > 0$ small enough that $\eta < s_\varepsilon$, every value in J is assigned when $\lambda_t \leq \eta$, and $\lambda_t \leq \eta$ implies $u_t < \bar{u} + \varepsilon$. At every such date after τ, the costate is positive. Let r_t be the unique unconstrained continuation, so $q\lambda_t \in \Lambda(r_t)$. Monotonicity and $q\lambda_t < \lambda_t$ give $r_t \leq u_t$. Therefore, when $V_t \in J$, the participation floor satisfies $\ell(V_t) > \bar{u} + \varepsilon > u_t \geq r_t$. Strict concavity above the plateau makes this floor the unique constrained continuation, so $u_{t+1} = \ell(V_t) \in O$.

Lemma 23, applied to η and J, shows that these dates have positive lower frequency, so O is reached infinitely often. Apply the argument to the countable collection of open intervals with rational endpoints whose closures lie in $(\bar{u}, u^\sharp)$. For any x in this interval, choose nested members of this collection that contain x and have diameters converging to zero. Successively later visits to these intervals form a subsequence converging to x. Thus every point of $(\bar{u}, u^\sharp)$ is a subsequential limit. The set of subsequential limits is closed, so it also contains both endpoints. The bounds $u_t \le u^\sharp$ at every date and $u_t \ge \bar{u}$ after τ give the reverse inclusion. Therefore $\omega(u_t) = [\bar{u}, u^\sharp]$ almost surely. This proves part (iii). □

A.6 Workforce Size and Task Pooling

This subsection proves Proposition 7. It first shows that, with one task per period, a vector of worker promises can be replaced by its sum. It then compares flexible matching with separate single-task relationships and constructs environments in which a worker beyond the number of tasks is strictly valuable.

Proof of Proposition 7. Part (i)*: one task.* Let $K = 1$. The feasible promise set is $\mathcal{U}^{N,1} = \{u \in \mathbb{R}^N_+ : \sum_i u_i \le V^+\}$. Put $s := \sum_i u_i$.

Take any feasible N-worker contract starting from u. Let X_t indicate whether the task is assigned at date t, and let $s_t := \sum_i u_{i,t}$. Summing promise keeping across workers gives

$$s_t \le \mathbb{E}_t\big[(1-\delta_A)V_t X_t + \delta_A s_{t+1}\big].$$

If the task is assigned to worker i, constraint (PC$_{ik}$) gives $u_{i,t+1} \ge \ell(V_t)$. Nonnegativity of the other promises therefore gives $s_{t+1} \ge \ell(V_t)X_t$. Thus (X_t, s_{t+1}) satisfies the one-worker promise-keeping constraint and (PC$_{ik}$) and gives the principal the same payoff. The constructed one-worker contract may condition its decisions on the full public history of the original contract.

At every such history, its current decision is feasible in the scalar Bellman problem. Iterating the Bellman inequality through date T gives

$$\Phi^{1,1}(s) \ge \mathbb{E}\left[(1-\delta_P)\sum_{t=0}^{T-1} \delta_P^t X_t + \delta_P^T \Phi^{1,1}(s_T)\right].$$

Letting $T \to \infty$ and using boundedness of $\Phi^{1,1}$ shows that the original contract's payoff is at most $\Phi^{1,1}(s)$. Hence $\Phi^{N,1}(u) \le \Phi^{1,1}(s)$.

For the reverse inequality, run an optimal scalar contract with worker i and leave every other worker idle. This gives $\Phi^{N,1}(se_i) \ge \Phi^{1,1}(s)$ for every i. If $s = 0$, the result follows. If $s > 0$, put $\alpha_i := u_i/s$. Since

$u = \sum_i \alpha_i s e_i$, concavity gives

$$\Phi^{N,1}(u) \geq \sum_i \alpha_i \Phi^{N,1}(se_i) \geq \Phi^{1,1}(s).$$

The corresponding public lottery selects worker i with probability α_i and runs the scalar contract with that worker. It delivers worker i expected utility at least $\alpha_i s = u_i$, which explains the concavity construction.

Part (ii)*: weak comparisons.* An optimal K-worker contract can be embedded in the N-worker problem by appending zero promises and leaving workers $K+1, \ldots, N$ idle. This gives $\Phi^{N,K}(0) \geq \Phi^{K,K}(0)$.

For the second inequality, the cube $[0, V^+]^K$ is contained in $\mathcal{U}^{K,K}$. For every $m \leq K$, the sum of the m largest nonnegative task values is at least m/K times the sum of all nonnegative task values. Taking expectations gives $\overline{V}_m \geq mV^+$. Hence, for every $x \in [0, V^+]^K$ and every worker set S of size m, $\sum_{i \in S} x_i \leq mV^+ \leq \overline{V}_m$, which proves the inclusion.

Now pair worker i permanently with task i and run K optimal scalar contracts. Each continuation coordinate remains in $[0, V^+]$, so the joint continuation vector lies in the feasible cube. The permanent pairs satisfy the matching constraints, and the total payoff is the sum of the scalar payoffs. Therefore $\Phi^{K,K}(0) \geq K\Phi^{1,1}(0)$.

Part (iii)*: strict value from flexible matching.* Suppose $K \geq 2$ and $\mathbb{E}[V] < \ell(\underline{v})$. Write $\phi := \Phi^{1,1}$, let $[0, \bar{u}]$ be its plateau, put $C := [0, V^+]^K$, and define $\Psi(x) := \sum_{r=1}^K \phi(x_r)$ for $x \in C$.

For $x \in C$, let $G(x)$ be the largest value of

$$\mathbb{E}[(1 - \delta_P)(\text{number of assignments}) + \delta_P \Psi(z)]$$

over relaxed one-period rules that match the current tasks, choose $z \in C$, satisfy (PC$_{ik}$) after every task vector, and deliver expected utility at least x. The set of relaxed rules is compact and convex, so the maximum is attained. Mixing feasible rules makes G concave, while lowering a coordinate of x enlarges the feasible set and makes G coordinatewise nonincreasing. Running the scalar Bellman-optimal rule on each fixed worker–task pair is feasible and attains $\Psi(x)$. Hence $G(x) \geq \Psi(x)$ throughout C.

Flexible matching strictly raises this one-period value whenever one promise is above the scalar plateau and another lies in its interior. To prove this, take an interior point $x \in C$ with $x_i > \bar{u}$ and $x_j \in (0, \bar{u})$, where $i \neq j$, and suppose that $G(x) = \Psi(x)$. Assigning each positive task to its fixed worker and setting every continuation promise equal to V^+ delivers $V^+ > x_r$ to every worker. This strict feasibility gives strong duality and a multiplier $\lambda \geq 0$ on promise keeping. The associated supporting inequality is $G(y) \leq G(x) - \lambda \cdot (y - x)$

for every $y \in C$. Because $G \geq \Psi$ and the functions agree at x, the same inequality supports $-\Psi$ at x. The scalar value is flat on the interior of its plateau and strictly decreasing above it, so $\lambda_i > 0$ and $\lambda_j = 0$.

Under the supposed equality, running the scalar optimal rule on each fixed pair attains $G(x)$. Strong duality implies that this rule maximizes the λ-Lagrangian. On the positive-probability event $0 < V_i < V_j$, both tasks are assigned to their fixed workers. Swap these two positive tasks and retain every continuation promise. The swap preserves matching feasibility, (PC_{ik}), the number of assignments, and the continuation payoff. It raises the Lagrangian by $(1-\delta_A)(\lambda_i - \lambda_j)(V_j - V_i) > 0$, a contradiction. Thus $G(x) > \Psi(x)$.

The proof of Proposition 4 identifies a common plateau assignment rule X_0 and shows that $E_0 := \{v : X_0(v) = 1,\ \ell(v) \in (\bar{u}, V^+]\}$ has positive probability. Since F has a density, the event $\{\ell(V) = V^+\}$ is null, so $E := E_0 \cap \{\ell(v) < V^+\}$ also has positive probability.

Choose $a \in (0, \bar{u})$. At scalar promise zero, use X_0. For an assigned task with $\ell(v) > \bar{u}$, set the continuation equal to $\ell(v)$. In every other case, set it to $\max\{a, \ell(v)X_0(v)\}$. The continuation lies on the plateau whenever a plateau promise satisfies (PC_{ik}) and otherwise equals the participation floor. It therefore maximizes the scalar continuation value conditional on $X_0(v)$. Each realization delivers nonnegative utility, so promise keeping at zero also holds.

Run this decision independently for the K fixed pairs in the first period. Independence, $F(E) > 0$, and $\Pr(V > 0) > 0$ imply that $\mathcal{B} := \{V_1 \in E,\ V_2 > 0, \ldots, V_K > 0\}$ has positive probability. On $\mathcal{B}$, the next promise vector is $x = (\ell(V_1), a, \ldots, a)$, with one coordinate in $(\bar{u}, V^+)$ and every other coordinate in $(0, \bar{u})$.

At the next date, on $\mathcal{B}$, select a measurable relaxed rule attaining $G(x)$; elsewhere begin the separate scalar continuations immediately. After the one-period maximizing rule, run separate scalar contracts from its continuation vector in C. Relative to separate operation, the date-zero gain is

$$\delta_P \mathbb{E}\big[\big(G(\ell(V_1), a, \ldots, a) - \Psi(\ell(V_1), a, \ldots, a)\big)\mathbf{1}_{\mathcal{B}}\big] > 0.$$

A maximizing relaxed rule can be selected measurably in x by the measurable maximum theorem, as in the proof of Lemma 2. Since C is compact and Ψ is concave and continuous, the argument in Lemma 3 applies: average each continuation lottery conditional on the matching, then purify the finite matching lottery. Concavity weakly raises the objective in the averaging step, and maximality of the relaxed rule makes this increase zero. Purification then preserves the principal's payoff and every worker's delivered utility. The parameterized selection gives a measurable deterministic rule in x. Thus the strict gain is attained by a deterministic recursive contract, proving part (iii).

Part (iv)*: strict value from an additional worker.* Write $\widetilde{\Phi}_F^{N,K}$ for the value when public randomization over contract decisions is allowed. Under every distribution satisfying the maintained density assumptions, Lemma 3 gives $\widetilde{\Phi}_F^{N,K} = \Phi_F^{N,K}$. We first construct an atomic three-point distribution F_0, for which the relaxed formulation is convenient.

Fix $K \geq 2$, set $\delta_A = 1/2$, and choose probabilities $g, m, e > 0$ with $g + m + e = 1$. A task is good, mild, or severe, with respective values 1, $-a$, and $-R_K$, where $R_K := 1 - (1 - g)^K$. Since $\delta_A = 1/2$, the participation floor is $\ell(v) = [-v]_+$. If $J \sim \text{Binomial}(K, g)$, the maximum total utility deliverable to any group of r workers is $b_r := \mathbb{E}[\min\{J, r\}]$ for $r = 1, \ldots, K + 1$. Thus $b_1 = R_K$, $b_{K+1} = Kg$, and R_K is the largest promise that one worker can receive. Moreover, $b_r > R_K$ for $r = 2, \ldots, K + 1$, while $2R_K - b_2 = \Pr(J = 1) > 0$.

Choose $\delta_P > 1/(1 + m^K)$, and let $\beta := 1 - (1 - e)^K$, the probability that at least one severe task arrives. Choose $a > 0$ sufficiently small that the following inequalities hold:

$$\begin{aligned} g - ma &> a, \\ ra &< b_r \quad (r = 1, \ldots, K + 1), \\ R_K + (r - 1)a &< b_r \quad (r = 2, \ldots, K + 1), \\ Kg &> a(Km + K + 1 + \beta), \\ Kg - R_K &> Ka(m + 1). \end{aligned}$$

All inequalities hold strictly at $a = 0$, so they continue to hold for all sufficiently small positive a. The bounds involving b_r ensure that the promise vectors below are feasible; the last two bounds ensure promise keeping.

First consider K workers. Reject every severe task, assign every good or mild task symmetrically, and give every worker continuation promise a. Each worker's current expected utility is $g - ma > a$. The vector $a\mathbf{1}_K$ is feasible because every r-coordinate sum is $ra < b_r$, and the inequality $g - ma > a$ gives promise keeping at this state. The policy is also feasible from promise zero and gives the principal $B := K(1 - e)$. We next show that this is the largest relaxed value with K workers. Let D_t be the number of nonsevere tasks at date t, and let I_t indicate that all K tasks are mild. Then $\mathbb{E}[D_t] = B$ and $\mathbb{E}[I_t] = m^K$.

Suppose a contract assigns a severe task. Constraint (PC$_{ik}$) gives its recipient a next-period promise of at least R_K, and the individual promise bound makes that promise exactly R_K. Put $M(v) := \max_k [v_k]_+$. A worker promised R_K satisfies

$$R_K \leq \mathbb{E}_t\left[\tfrac{1}{2}Y_i + \tfrac{1}{2}u_i'\right] \leq \tfrac{1}{2}\mathbb{E}[M(V)] + \tfrac{1}{2}R_K = R_K.$$

The two pointwise gaps $M(V) - Y_i$ and $R_K - u_i'$ are nonnegative, so equality of expectations makes both zero almost surely. Thus this worker must receive a good task whenever one is available, remain idle when all tasks are negative, and retain promise R_K. She cannot receive another severe task. A severe task also cannot be assigned to another worker, because the resulting two continuation promises of R_K would violate $2R_K > b_2$.

Let τ be the first date on which a severe task is assigned. Before τ, the contract completes at most D_t tasks. At τ, at most one severe task can be assigned, so it completes at most $D_\tau + 1$. After τ, no severe task can be assigned. On an all-mild draw, the worker with promise R_K must remain idle, so at most $D_t - I_t$ tasks are completed. Conditional on the history through τ, independence of later draws gives

$$(1-\delta_P)\mathbb{E}\left[\sum_{t=\tau+1}^{\infty} \delta_P^t I_t \,\middle|\, \text{history through } \tau\right] = \delta_P^{\tau+1} m^K.$$

Comparing with the policy that always completes every nonsevere task gives, with the term set to zero when $\tau = \infty$,

$$\Pi \le B + \mathbb{E}\big[\delta_P^\tau\big((1-\delta_P) - \delta_P m^K\big)\big] \le B.$$

The term $1 - \delta_P$ is the largest current gain from the first severe assignment, and $\delta_P m^K$ is the ensuing discounted loss. The choice of δ_P makes the bracket negative. Hence $\widetilde{\Phi}_{F_0}^{K,K}(0) = B$.

Now give the firm $K+1$ workers. Define $s := a\mathbf{1}_{K+1}$ and, for each worker i, $r_i := R_K e_i + a(\mathbf{1}_{K+1} - e_i)$. For s, every r-coordinate sum is $ra < b_r$. For r_i, a set of r coordinates excluding i sums to $ra < b_r$; a set containing i sums to $R_K = b_1$ when $r = 1$, and to $R_K + (r-1)a < b_r$ when $r \ge 2$. Thus all these states belong to $\mathcal{U}^{K+1,K}$.

At s, assign every nonsevere task and one severe task whenever a severe task appears. Randomize symmetrically over the entire matching. If every task is nonsevere, remain at s; otherwise continue at r_i, where i is the symmetrically selected severe-task recipient. At r_i, give worker i a good task whenever one is available. Allocate all remaining nonsevere tasks symmetrically among the other K workers, reject severe tasks, and remain at r_i.

At r_i, worker i's expected current utility and continuation promise both equal R_K, so promise keeping binds. Each other worker has expected current utility $(Kg - Kma - R_K)/K$ and continuation a. Their constraints are satisfied because $Kg - R_K > Ka(m+1)$. At s, symmetry makes the coordinatewise constraints equivalent to their sum. Expected current utility summed over workers is $K(g - ma) - \beta R_K$, and the expected sum of next promises is $(1-\beta)(K+1)a + \beta(R_K + Ka)$. The sum of current and continuation utility, each weighted by one half, is at least $(K+1)a$ exactly when $Kg \ge a(Km + K + 1 + \beta)$. Constraint ($\mathrm{PC}_{ik}$) holds as well: at s, a mild-task recipient continues at $a = \ell(-a)$ and the severe-task recipient at $R_K = \ell(-R_K)$; at r_i, only good and mild tasks are assigned, and each mild-task recipient continues at a. Thus the policy is feasible.

The stationary policy at each r_i has payoff B. If W is the payoff from state s, then

$$W = (1-\delta_P)(B+\beta) + \delta_P\big[(1-\beta)W + \beta B\big], \qquad W - B = \frac{(1-\delta_P)\beta}{1-\delta_P(1-\beta)} > 0.$$

The policy is also feasible from initial promise zero. Therefore $\widetilde{\Phi}^{K+1,K}_{F_0}(0) \geq W > B = \widetilde{\Phi}^{K,K}_{F_0}(0)$.

It remains to approximate F_0 by distributions satisfying the maintained density assumptions. The severe value reaches the individual promise bound under F_0, so we place the approximating severe values slightly above that bound. Work on the common support $[-1,1]$. Choose sequences $\varepsilon_n \downarrow 0$ and $\eta_n \downarrow 0$. Give total mass $1-\eta_n$, in proportions approaching g, m, e, to three disjoint smooth bumps, and give mass η_n to a uniform component. Put the good bump in $(1-2\varepsilon_n, 1-\varepsilon_n)$ and the mild bump in $(-a+\varepsilon_n, -a+2\varepsilon_n)$. The uniform component makes each density strictly positive and bounded on the full support.

Let $R_{K,n} := \overline{V}_1(F_n)$. This number depends only on the positive part of F_n, so it can be determined before locating the severe bump. Put that bump in $(-R_{K,n}+\varepsilon_n, -R_{K,n}+2\varepsilon_n)$. Then every mild-bump value exceeds $-a$, and every severe-bump value exceeds $-R_{K,n}$. Their participation floors are therefore below a and $R_{K,n}$, respectively. The distributions F_n converge weakly to F_0, and $\overline{V}_r(F_n) \to b_r$ for every r.

For the $K+1$-worker policy, retain $s = a\mathbf{1}_{K+1}$ and replace r_i by $r_{i,n} := R_{K,n}e_i + a(\mathbf{1}_{K+1} - e_i)$. The strict subset inequalities for s and for every subset of size at least two in $r_{i,n}$ continue to hold for large n; the remaining coordinate bound is $R_{K,n} = \overline{V}_1(F_n)$. Hence these states are feasible.

At s, use the preceding symmetric rule when every draw belongs to one of the three main bump regions. If any draw lies outside those regions, assign nothing and remain at s. At $r_{i,n}$, give worker i the greatest positive draw, allocate the remaining positive draws and mild-bump draws symmetrically among the other workers, reject all other negative draws, and remain at $r_{i,n}$. Worker i's expected current utility is exactly $\mathbb{E}[\max_k [V_k]_+] = R_{K,n}$, equal to her continuation promise. The other workers' expected utilities and the aggregate utility delivered at s converge to the strict bounds verified under F_0, so their promise-keeping inequalities hold for all large n. The placement of the mild and severe bumps ensures (PC_{ik}) directly.

Let B_n be the value of this policy at any state $r_{i,n}$, and let W_n be its value at s. The policy has finitely many states, and its flow payoffs and transition probabilities converge to those under F_0. Thus $B_n \to B$ and $W_n \to W$. The same first-period rule is feasible from promise zero and has payoff W_n, which gives

$$\liminf_{n\to\infty} \widetilde{\Phi}^{K+1,K}_{F_n}(0) \geq W.$$

We next establish the required upper bound for K workers. Let $H^{K,K}_{F,T}(u)$ be the largest relaxed payoff over the first T periods from promise u, with zero terminal principal payoff and terminal promises restricted to $\mathcal{U}_F^{K,K}$. Every feasible terminal promise can be delivered by a continuation contract, while the principal's payoff after date T is at most $K\delta_P^T$. Therefore

$$H_{F,T}^{K,K}(0) \leq \widetilde{\Phi}_F^{K,K}(0) \leq H_{F,T}^{K,K}(0) + K\delta_P^T.$$

For each fixed T, we claim more generally that if $u_n \in \mathcal{U}_{F_n}^{K,K}$, $u_n \to u$, and $u \in \mathcal{U}_{F_0}^{K,K}$, then

$$\limsup_{n\to\infty} H_{F_n,T}^{K,K}(u_n) \leq H_{F_0,T}^{K,K}(u).$$

The claim is immediate for $T = 0$. For the induction step, take nearly optimal relaxed one-period laws. Task vectors, matchings, and continuation promises lie in a common compact space, so a subsequence converges weakly. Its task marginal is F_0^K. Convergence of the bounds $\overline{V}_r(F_n)$ gives Hausdorff convergence of the promise polytopes. The matching and promise-set restrictions and (PC_{ik}) are closed, and bounded current utilities make promise keeping pass to the limit. The weak limit is therefore a feasible relaxed decision under F_0. The induction hypothesis gives

$$\limsup_{n\to\infty} H_{F_n,T-1}^{K,K}(z_n) \leq H_{F_0,T-1}^{K,K}(z) \quad \text{whenever } z_n \to z,$$

and these value functions are uniformly bounded by K. The extended Portmanteau theorem for varying integrands therefore bounds the expected continuation payoff under the convergent one-period laws and gives the displayed inequality.

At $u = 0$, the finite-horizon inequality and the preceding payoff bound give

$$\limsup_{n\to\infty} \widetilde{\Phi}_{F_n}^{K,K}(0) \leq H_{F_0,T}^{K,K}(0) + K\delta_P^T \leq B + K\delta_P^T.$$

Letting $T \to \infty$ yields $\limsup_n \widetilde{\Phi}_{F_n}^{K,K}(0) \leq B$. Since $W > B$, every sufficiently large n satisfies

$$\widetilde{\Phi}_{F_n}^{K+1,K}(0) > \widetilde{\Phi}_{F_n}^{K,K}(0).$$

Each F_n has a continuous, strictly positive, bounded density. At every promise state, Lemma 3 preserves the one-period principal payoff, including continuation value, and all delivered-utility coordinates while replacing the relaxed matching rule by a measurable deterministic rule. Hence $\widetilde{\Phi}_{F_n}^{N,K} = \Phi_{F_n}^{N,K}$, and the last strict inequality proves part (iv). □

References

Abdulkadiroğlu, A., & Bagwell, K. (2013). Trust, reciprocity, and favors in cooperative relationships. *American Economic Journal: Microeconomics*, *5*(2), 213–259.

Abreu, D., Pearce, D., & Stacchetti, E. (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, *58*(5), 1041–1063.

Akerlof, G. A., & Yellen, J. L. (1990). The fair wage–effort hypothesis and unemployment. *The Quarterly Journal of Economics*, *105*(2), 255–283.

Andrews, I., & Barron, D. (2016). The allocation of future business: Dynamic relational contracts with multiple agents. *American Economic Review*, *106*(9), 2742–2759.

Baker, G., Gibbons, R., & Murphy, K. J. (1994). Subjective performance measures in optimal incentive contracts. *The Quarterly Journal of Economics*, *109*(4), 1125–1156.

Balseiro, S. R., Gurkan, H., & Sun, P. (2019). Multiagent mechanism design without money. *Operations Research*, *67*(5), 1417–1436.

Bewley, T. F. (1999). *Why wages don't fall during a recession*. Cambridge, MA: Harvard University Press.

Bird, D., & Frug, A. (2019). Dynamic non-monetary incentives. *American Economic Journal: Microeconomics*, *11*(4), 111–150.

Bird, D., & Frug, A. (2021). Optimal contracts with randomly arriving tasks. *The Economic Journal*, *131*(637), 1905–1918.

Bird, D., & Frug, A. (2022). Monotone contracts. *Theoretical Economics*, *17*(3), 1041–1073.

Casella, A. (2005). Storable votes. *Games and Economic Behavior*, *51*(2), 391–419.

Castro, F., Ma, H., Nazerzadeh, H., & Yan, C. (2021). Randomized FIFO mechanisms. *arXiv preprint arXiv:2111.10706*.

Combe, J., Nora, V., & Tercieux, O. (2025). Dynamic assignment without money: Optimality of spot mechanisms. *Theoretical Economics*, *20*(1), 255–301.

de Clippel, G., Eliaz, K., Fershtman, D., & Rozen, K. (2021). On selecting the right agent. *Theoretical Economics*, *16*(2), 381–402.

Dvoretzky, A., Wald, A., & Wolfowitz, J. (1951). Elimination of randomization in certain statistical decision procedures and zero-sum two-person games. *The Annals of Mathematical Statistics*, *22*(1), 1–21.

Eliaz, K., Fershtman, D., & Frug, A. (2025, August). *Clerks* (Working Paper No. 1535). Barcelona School of Economics.

Forand, J.-G., & Zápal, J. (2020). Production priorities in dynamic relationships. *Theoretical Economics*, *15*(3), 861–889.

Frankel, A. (2016). Discounted quotas. *Journal of Economic Theory*, *166*, 396–444.

Fudenberg, D., & Rayo, L. (2019). Training and effort dynamics in apprenticeship. *American Economic Review*, *109*(11), 3780–3812.

Goldlücke, S., & Tröger, T. (2018). Assigning an unpleasant task without payment. *Collaborative Research Center Transregio 224 Discussion Paper Series*, *3*.

Guo, Y., & Hörner, J. (2018). Dynamic allocation without money. *Working Paper*.

Holmström, B. (1984). On the theory of delegation. In M. Boyer & R. E. Kihlstrom (Eds.), *Bayesian models in economic theory* (pp. 115–141). Amsterdam: North-Holland.

Hylland, A., & Zeckhauser, R. (1979). The efficient allocation of individuals to positions. *Journal of Political Economy*, *87*(2), 293–314.

Jackson, M. O., & Sonnenschein, H. F. (2007). Overcoming incentive constraints by linking decisions. *Econometrica*, *75*(1), 241–257.

Ji, Y., & Vong, A. (2025). Task assignment as dynamic incentives. *Working Paper*. (arXiv:2511.05338)

Krasikov, I., Lamba, R., & Mettral, T. (2023). Implications of unequal discounting in dynamic contracting. *American Economic Journal: Microeconomics*, *15*(1), 638–692.

Marshall, A. W., Olkin, I., & Arnold, B. C. (2011). *Inequalities: Theory of majorization and its applications* (2nd ed.). New York: Springer.

Myerson, R. B. (1982). Optimal coordination mechanisms in generalized principal–agent problems. *Journal of Mathematical Economics*, *10*(1), 67–81.

Prescott, E. S., & Townsend, R. M. (2006). Private information and intertemporal job assignments. *The Review of Economic Studies*, *73*(2), 531–548.

Samuelson, L., & Stacchetti, E. (2017). Even up: Maintaining relationships. *Journal of Economic Theory*, *169*, 170–217.

Spear, S. E., & Srivastava, S. (1987). On repeated moral hazard with discounting. *The Review of Economic Studies*, *54*(4), 599–617.