SUBMITTING TO AUTHORITY:

AGREEING TO FOLLOW EX ANTE UNKNOWN ORDERS

August 23, 2025

Birger Wernerfelt*

Abstract

We show that contracts under which workers agree to follow ex ante unknown orders in return for fixed hourly wages can be asymptotically efficient. Players do not need to know the set of possible orders nor ever observe each other's payoffs; they just need to know the distributions of costs and values. This allows them to sustain a relational contract in the form of a quota mechanism even if costs have infinite support. The mechanism lets the players identify inefficient orders such that they can avoid implementing them. The prescribed behavior shares many properties with employment and the conditions under which it is efficient, frequent and diverse changes, are similar to those under which employment is used. (*JEL* D01, D02, D86, J41, L14) (*Key Words:* Authority, Incomplete contracts, Employment, Quota mechanisms.)

In a recent job posting MIT was looking for a salaried academic secretary to

- perform a variety of administrative tasks for MIT.nano,

- assist with conference and event planning and coordinating,

- provide complex administrative office support,

- coordinate conference rooms bookings,

- maintain and update portions of facility website content,

- act as the primary point of contact and customer success representative, and

- support existing users with administrative inquiries, registration, and onboarding activities.

* bwerner@mit.edu, I am grateful to Ian Ball for several extremely helpful comments. Any errors are, of course, the responsibility of the author.

The labor contract offered in this ad is very different from those we usually write about in pages like these. The job description includes seven categories of tasks each of which has a continuum of elements, costs vary wildly, you never observe the opponent's payoff, and per-period pay is constant. It is also an example of a contract that is very common. Why?

Noting that task-by-task negotiation would be prohibitive and that complete contracting is impossible, we here identify a mechanism which overcomes the above-mentioned complications and asymptotically implements the first best in a perfect equilibrium. The model does not assume that the players have any ex ante knowledge of the set of possible tasks, only that they know the probability distributions of their costs and values. Once a specific task is identified at the start of a period, the boss learns its value and the worker learns its cost, but that does not teach them anything about any other tasks. There are some tasks for which cost exceed value and in spite of the fact that payoffs are ex post unobservable, the players can use quotas to identify and eliminate such tasks. Specifically, the worker may make an unverifiable claim to the effect that a task has high costs. If such a claim is made, the boss may accept it and withdraw the order but could also make an unverifiable counterclaim to the effect that the task has high value, implying that it should be performed in spite of the high cost. Using their knowledge of the value and cost distributions, players can assess the veracity of the claims by comparing the actual frequency of high cost (value) claims with theoretical frequencies over large numbers of tasks. To ensure that players prefer truth-telling over frontloading of claims, the result requires that orders are issued frequently such that the inter-task discount rate is very small.[1]

There is a small but recently growing theoretical literature on such quota mechanisms. The first example was probably Radner's (1985) "review" strategies in a game in which two players sequentially and repeatedly report their private information (he looked at a Principal-Agent model in which the principal did statistical tests on the agent's output.) More recently Jackson and Sonnenschein (2007) and Ball, Jackson, and Kattwinkel (2022) derived some general efficiency results in games in which information arrives simultaneously, while Renault, Solan,

[1] The critical role played by the law of large numbers is consistent with the stylized fact that this type of contract is not used for individual tasks, for tasks that take a long time to complete, or when new tasks only are needed infrequently. There is very little systematic evidence about the frequency with which orders are given, but the data we do have are quite compelling. In particular, Guest (1968) followed 56 foremen in an auto plant and found that they attended to a new issue every 48 seconds. They did not give orders that frequently but still spent an average of 26.4% of their time face to face with direct reports.

and Vieille (2013), Escobar and Toikka (2013), Frankel (2016), and Ball and Kattwinkel (2025, Section 8.2) consider dynamic games such as ours. The case we consider is arguably closest to that of Escobar and Toikka except that our players move sequentially and one of them (the worker) has an infinite type space.[2] To the best of our knowledge this is the first paper on quota mechanisms in which one player has infinite type space.

In spite of the widespread use of incomplete contracts with submission to authority and fixed payments, we still do not have models that reflect the full complexity of order taking. The first paper on the topic is probably Simon (1951), who suggests that workers take orders in return for fixed payments when they are roughly indifferent between the tasks covered by the agreement. In the present paper, the costs can have a very large variance as long as their distribution is known. Another early contribution is Alchian and Demsetz (1972), who use an analogy to a customer buying products from a grocer. The latter situation differs from that considered here for two reasons: Each product has its own price and the set of products is finite and known ex ante. In the more recent literature several authors, e. g. Baker, Gibbons, and Murphy (2002), have like us looked at the employment relationship through the lens of relational contracting but make stronger informational assumptions and often focus on moral hazard. The work closest to ours is Bolton and Dewatripont (2013, Ch. 12.5) who also allow payoffs to be unobservable but assume that the boss has superior ex ante information about the worker's costs and rely on bonuses to achieve efficiency.

A different and very influential approach is pioneered by Grossman and Hart (1986) and Hart and Moore (1990), who assume that contracts necessarily are incomplete but allow ex post bargaining (renegotiation) such that pay can vary from task to task. Another stream, prominently associated with Williamson, assumes that agents are boundedly rational and that contracts necessarily are incomplete. He (1975, p. 66) explains the prevalence of simple contracts by the possibility that more complicated contracts could be exploited by opportunistic agents.[3] In

[2] In our application it is clearly a realistic assumption that costs have infinite support and our treatment of that case constitutes a small theoretical contribution. The existence proof relies on the fact that the player with finite type space has a dominant strategy such that the opponent can be assumed to maximize against that. We are then able to characterize that maximum in enough detail to derive a limiting efficiency result as the discount rate goes to zero.

[3] Consistent with this theory, Levin and Tadelis (2010) argue that a contract on time is easier to administer than a contract on quality and use this to explain several empirical regularities.

contrast to this, the contract in the present paper is asymptotically optimal while still being very simple.

The basic model and the main results are presented in Section I, a few extensions are sketched in Section II, and a discussion concludes the paper in Section III.

I. MODEL

Time is discrete and infinite and one of an infinite number of tasks is randomly chosen at the start of each period. The task can be performed during that period but not later. The boss (she) derives value from having tasks performed and the worker (he) can do it but incurs costs in the process. The value of task i is v_i where $v_i \in \{v, V\}$ and its cost is c_i where c_i is drawn from the distribution F with support $(0,V)$. We define $c \equiv \int_0^v xf(x)dx/F(v)$ and $C \equiv \int_v^V xf(x)dx/[1 - F(v)]$, such that c is the mean cost if costs are below v, and C is the mean cost if costs are above v. The probability of $c_i > v$ is $1 - F(v) \equiv p$ and the probability of $v_i = v$ is q. Values and costs are independently drawn and trade is inefficient iff $c_i > v$ and $v_i = V$ such that the first best expected per-period surplus is $(v - c)(1- p)q + (V - C)p(1 - q) + (V - c)(1 - p)(1 - q) \equiv \Gamma^*$. Value and cost realizations are, even ex post, known only to the boss and worker respectively, although F, v, V, and q are common knowledge. There is transferable utility, the players are risk neutral and discount future payments at the same rate $r \geq 0$ per period.

We will be looking for mechanisms that asymptotically can implement the first best outcome. One possibility is that the parties engage in independent negotiations every period. It is reasonable to assume that they bargain efficiently but it also reasonable to assume that the individual negotiations are not completely costless (because of the associated delays and other costs). These costs will be unimportant in many settings particularly if costs and values are high and periods are long. They will, however, cumulate to a non-trivial factor if individual trades are small and frequent as is often the case for trades in human asset services. In such circumstances, as when an executive needs a stream of different services from her secretary, we will assume that the surplus achievable with trade-by-trade negotiation is less than Γ^*.

Alternatively, the traders could try to write a complete contingent long-term contract, but since the number of possible tasks is infinite, this cannot be done.

We will, however, show that they can agree to a relational contract that exhibits the main characteristics of the employment contracts such that (1) the boss gives an order every time a new task is needed, (2) the worker is paid an ex ante agreed-upon constant wage in every period, (3) in the limit, the worker obeys the order unless trade is inefficient, and (4) the contract is at will.

The idea is to use *quota strategies* as a point of reference. These strategies are defined on blocks of tasks (or equivalently, periods) as follows: The worker can claim that cost is "high" up to a fixed number of times per block. If the worker makes a claim, the boss can respond with the counterclaim that value is "high" up to another fixed number of times per block, and if either trader exceeds their quota the parties terminate their relationship.

With the restriction to quota strategies we assume that the traders, after an initial once-and-for-all agreement on a constant per-period price, play the following stage game each period:

1. The boss learns the identity and value of the next task.
2. She orders the worker to perform it and he learns its cost.
3. The worker may claim that the costs are high by announcing C.
4. If he does, the boss may make the counterclaim that the value is high by announcing V.
5. The worker may or may not perform the task.
6. The boss pays the agreed-upon per-period price.
7. If either player exceeds their quota in a block, the relationship ends, and the players revert to task-by-task negotiation.
8. After b periods, a new block commences and the counting of claims and counterclaims restarts from zero.

We aim to show the following:

Proposition: *If $\Gamma(e)$ is the expected per-period surplus in the perfect equilibrium e,*
$\forall\, \varepsilon > 0\ \exists\, b' > 0\ \forall\, b > b'\ \exists\, r' > 0\ \forall\, r < r'\ \exists\, e\colon \Gamma(e) + \varepsilon > \Gamma^*$

It is helpful to start with the non-equilibrium *simple quota strategies*:

(a) If possible, the worker claims high cost the first *pb* times $c_i > v$. If he, in *the end of a block,* has *k* claims left when *k* tasks remain in the block, he ends by making *k* straight claims for a total of *pb;* regardless of his actual costs in those periods. (So workers use up their quotas exactly; in some cases by making false claims and in others by not making true claims, both at the very end of a block.)

(b) If possible, the boss responds to a worker's claim by making the counterclaim that value is high whenever $v_i = V$. If she, in *the end of a block*, has *K* counterclaims left when the worker has *K* claims left, she ends by making *K* straight counterclaims for a total of $(1 - q)pb$; regardless of her values in those periods. (Bosses know that they will get *pb* chances to counterclaim and they will use up their quotas exactly; in some cases by making false counterclaims and in others by not making true counterclaims.)

(c) The worker follows the order unless he has claimed high cost and the boss has not made a counterclaim.

The proof works by first showing that the simple quota strategies asymptotically implement the first best as blocks get longer and discount rates get smaller (because the frequencies with which high cost and value occur converge to *p* and $1 - q$, respectively). We next characterize a pair of *equilibrium quota strategies* and show that their expected per-period surplus converges to that from the simple quota strategies, and thus the first best, as blocks get longer and discount rates get smaller.

Lemma 1: *If* $\Gamma(e_p)$ *is the expected per-period surplus when both players follow the simple quota strategies,* $\forall\, \varepsilon > 0\ \exists\, b' > 0\ \forall\, b > b'\ \exists\, r' > 0\ \forall\, r < r' : \Gamma(e_p) + \varepsilon > \Gamma^*$

Proof: If in each block there are exactly *pb* realizations where costs are greater than *v* and exactly $(1-q)pb$ of these where per-period values are *V*, then the expected surplus created by the above strategies are $(v - c)(1- p)q + (V - C)p(1 - q) + (V - c)(1 - p)(1 - q) = \Gamma^*$ per-period. Since the cost and value realizations are random, there will in many blocks be deviations from the mean and consequently some end-of-block inefficiency. As described in the previous

paragraph, these come in four variants: (i) If the worker has to make false claims, some trades with expected gains v-c will be missed, (ii) if the worker cannot make enough claims, some trades with expected gains v-C will be made, (iii) if the boss has to make false claims, some trades with expected gains v-C will be made, and (iv) if the boss cannot make enough claims, some trades with expected gains V-C will be missed. As is obvious, only one of (i) and (ii) and only one of (iii) and (iv) can happen in any one block. The expected per-period loss from these inefficiencies depends on $F, v, V, q, r,$ and b. In particular, since the standard error of a sum of b binomial variables is proportional to $b^{1/2}$, the expected error per-period is proportional to $b^{-1/2}$ and we can bring the expected undiscounted per-period loss below any ε simply by choosing a sufficiently large b.

QED

Remembering that the inefficiencies are caused by histories with surprisingly many or surprisingly few high cost or high value realizations, we now proceed to show that there exists a perfect equilibrium in which expected per-period surplus converges to the first best as blocks get longer and discount rates get smaller.

Lemma 2: *In any equilibrium, both players will exhaust their quotas in each block.*

Proof: The worker clearly prefers to make as many claims as allowed since each claim gives him a chance of getting paid in spite of incurring zero costs. Similarly, the boss prefers to counterclaim high value (and get a task performed) if the value in fact is high, but even if the value is low, it is still more attractive to counterclaim than not.

QED

Lemma 3: *For sufficiently low r, the simple quota strategy is dominant for the boss*.

Proof: Given a value of $b,$ this condition is satisfied if two things are true. First, if possible, she will not refrain from making a true counterclaim. The boss faces the strongest such temptation if the very first task is (high cost, V), the worker makes a claim, and she believes that he will react to her making a counterclaim by claiming high costs in each of the next $pb - 1$ straight tasks - the harshest punishment available to him. It is most attractive to stay quiet if she further believes that any further counterclaims will go unpunished. In that case, if she does make the counterclaim,

she gets V but then faces $pb - 1$ tasks with an average of $q0 + (1 - q)V$ followed by $(1 - p)b$ realizations with average of $qv + (1 - q)V$. So r' has to be sufficiently small to ensure that getting the early V is worth more than the expected cost of delaying the $(1 - p)qbv$. Second, r also has to be so small that it will also be a bad idea to spend a counterclaim on a v realization before the end-of-the-block. While doing so could give the boss a v (rather than a 0) earlier, it comes at the cost of a later V with probability $1 - q$. If r meets these two conditions, the simple quota strategy is dominant for the boss and the worker knows that she will play it.

QED

This means that the existence of equilibrium is not an issue: The worker can maximize against the simple quota strategy and since that maximization problem is bounded, it has a solution. Let us now characterize the worker's equilibrium strategy and its implications for the expected surplus.

Lemma 4: *Suppose that $r = 0$, that the worker just learned the cost of a new task, that $n < b$ tasks remain in the block, and that he has $k \leq pb$ claims left. As $b \rightarrow \infty$, the worker's reservation value for using one of those claims now goes to the expected cost of the k'th most costly task among the n yet to be revealed.. If $r > 0$, the reservation value will be lower.*

Proof: If $r = 0$ and all claims had the same chance of being met by a counterclaim, the behavior suggested by the Lemma maximizes his expected payoffs. Unfortunately, the second premise is only approximately true. The problem is that the boss, at the end-of-the-block, may have run out of counterclaims or may have "too many" left such that she will counter every claim even if her value is v. As $b \rightarrow \infty$ these possibilities weigh less in the worker's calculations and the postulated decision rule approximates equilibrium (optimal) behavior more closely. If $r > 0$, the future costs should be discounted such that the current nominal reservation value should be lower.

QED

Before getting to Lemma 5, it helps to look at an example in which we use the standard result (from the theory of order statistics[4]) that the mean of the k'th highest element in a sample of size n from $U[0, 1]$ is $(n + 1 - k)/(n + 1)$.

[4] Se e.g. David and Nagaraja, 2004.

Example. Suppose that $r = 0$, $F() = U[0, 1]$, and that $v = 2/3$ (such that $p = 1/3$). Suppose that the worker has one claim left $(k = 1)$. If one further task remains in the block $(n = 1)$, he will make a claim now if his costs are above $1/2$ - the expected cost of the single task still to be revealed. If two further tasks remain ($n = 2$), the expected maximum cost of the last two tasks and thus the reservation value is exactly v (= 2/3) and if three tasks remain it is $3/4$. At the start of the block $n = b - 1$, $k = pb$, and the reservation value is $(b - 1 + 1 - pb)/(b - 1 + 1) = 1 - p = v$, as it should be.

Lemma 5: *Compared to that resulting from the simple quota strategy, the equilibrium distribution of claim times differs in two ways: The expected end-of-block inefficiencies are fewer and if $r > 0$ the distribution is shifted to the left.*

Proof: The simple quota strategy requires that the worker, except for end-of-block effects, make claims iff his costs exceed v. However, as we saw in the example, he can do better by making claim decisions as functions of the number of remaining claims and tasks. Equilibrium reservation values typically differ over time as the worker reacts to the evolving history of high-cost tasks. If he has made fewer (more) claims than expected, he lowers (raises) the expected reservation cost below (above) v and thereby increases (slows) the expected speed, leading to fewer expected end-of-block inefficiencies. Additionally, if $r > 0$, the worker prefers to get his breaks earlier and equilibrium reservation values will be lower than if $r = 0$, causing the distribution of claim times to shift to the left.

QED

The reduction in expected end-of-block inefficiencies helps the boss as well, while the left skew hurts her. We are now ready to show:

Proposition 1: $\forall\, \varepsilon > 0\ \exists\, b' > 0\ \forall\, b > b'\ \exists\, r' > 0\ \forall\, r < r'\ \exists\, e'\colon \Gamma(e') + \varepsilon > \Gamma^*$

Proof: Comparing equilibrium payoffs to those when both play their simple quota strategies, the worker does better by definition. The comparison is more complicated for the boss: She will lose exactly $p(1-q)b$ tasks in both cases but the difference between the worker's strategy and the simple quota strategy means that her losses will differ in two ways: She will face fewer end-of-block inefficiencies and the worker will make claims earlier in the blocks. As $r \to 0$, the mutually beneficial reduction in end-of-block inefficiencies remains, but the left shift of the

claim distribution vanishes and with it the negative effect on the boss - who at the same time will care less and less about it. So her expected per-period equilibrium payoffs converge towards those when both parties use simple quota strategies. Recalling that the worker does better in equilibrium than if both use simple quota strategies, we see that their combined expected per-period surplus converges towards the first best as $r \to 0$.

QED

II. EXTENSIONS

The model is obviously very simple and we will now suggest several ways in which it can be generalized to fit more stylized facts.

Periods with no work

On some such occasions, the boss has no immediate needs, although she surely will have some in future periods. Specifically, suppose that the value of task i is v_i where $v_i \in \{0, v, V\}$ and that its cost is c_i where $c_i \in \{c_o, c, C\}$. When the boss has no needs, which happens with probability z, values and costs are 0 and c_o, respectively. In all other cases, the probability of C is p while that of the independently drawn V is q. We assume that it would be costly for the traders to find alternative partners, such that it is efficient for them to continue the relationship through periods in which the boss has no needs. In other periods, we look at the case in which $0 < c < v < C < V$, implying that trade is efficient unless $c_i = C$ and $v_i = v$ and the maximum achievable per-period surplus is $(1 - z)[(v - c)(1 - p)q + (V - C)p(1 - q) + (V - c)(1 - p)(1 - q)] - zc_o$. We can now proceed as in the main model.

The tasks differ in their duration

Suppose that the duration of task i, call t_i, is common knowledge once i is identified. Assume further that t_i is drawn independently of v_i and c_i from a distribution with mean τ and support $(0, T)$. In this case we weigh the claims and counterclaims by the durations of the associated tasks, define the quotas on the weighted claims and counterclaims, and proceed as in the main model. The biggest complication is that the reservation values in the worker's equilibrium strategy now are defined on $c_i t_i$.

Correlation between duration and per-period cost

Suppose that the duration of a random task is T or t with probabilities s and $1 – s$, respectively and that the probability of $c_i = C$ is P in the former case and p in the latter. In that case a similar equilibrium exists except that the quotas depend on the commonly known duration such that the worker can claim sPb times in a block if the duration is T and $(1 – s)pb$ times if it is not. The boss can respond to these claims by announcing V a total of $qsPb$ and $q(1 – s)pb$ times, respectively. We can now fill in the details as in the main model.

Granting the worker more time

Consider a task i with per-period value and cost (v, C) and suppose that the boss can reduce the worker's per-period cost to c' by granting him $1 + \Delta$ periods instead of just one to complete the task. Specifically, we assume that the total value remains v, but that the cost goes to $(1 + \Delta)c'$. By analogy to the equilibrium looked at in the main model, the boss allows the worker to claim C a total of pb times per block, and the worker in turn allows the boss to counterclaim V a total of $p(1 – q)b$ times per block. If the boss does not claim V, she gives the worker an additional Δ periods to complete the task and he accepts that offer. This may well be more efficient than dropping the trade: When inefficient trades are dropped, the maximum per-period surplus is Γ^*. To find the same number when workers may be granted more time, imagine that we pick a random period. The probability that value and costs are either $(V - C)$, $(V - c)$, or $(v - c)$ is proportional to $1 - pq$ and the probability that the period is one in which the worker works on a task for which he has been granted more time is proportional to $pq(1 + \Delta)$. So the probabilities that our random period is of the former or latter kind are $[1 - pq]/[1+pq\Delta]$ and $pq(1 + \Delta)/[1+pq\Delta]$, respectively. Since the value created in the two cases are Γ^* and $[v/(1 + \Delta) -c']$, the maximum surplus when the boss can grant more time is $\{[1 - pq]\Gamma^*+ pq(1+\Delta)[v/(1+\Delta) - c']\}/[1+pq\Delta]$ and this is greater than Γ^* iff $v/(1 + \Delta) – c' > \Gamma^*$. Intuitively, the left side is the per-period value if the worker is given extra time and the right side is the opportunity cost of that time. So small values of c', Δ, and Γ^*, as well as large values of v makes it more attractive to grant extra time.

The value distribution is multinomial.

The result continues to hold but the messages need to be richer and the blocks longer. For example, if values can be v_l , v_m, and v_h, the worker's messages need to specify whether costs are between v_l and v_m or higher than v_m and the buyer's messages need to differentiate between v_m and v_h. Since a smaller fraction of tasks on the average falls in each pair of realizations, per-cell convergence will be slower and a larger block size is needed.[5]

III. DISCUSSION

The mechanism shares many properties with typical employment contracts:

-The per-period wage is agreed upon before any work is done and the agreement includes a job description that helps the worker estimate the difficulty of the job as a whole, here captured by F.

-New orders are given very frequently.

-The orders are generally obeyed without new negotiations.

-Some individual tasks may not be worth the wage, but the average completed task is.

-Too much complaining, too many hard tasks, or refusal to follow an order will lead one of the parties to terminate the contract.

-The contract is open ended.

The environment depicted in the model is similar to those in which employment contracts are used. In particular, new tasks are needed very frequently, are drawn from an extremely large set, and vary in difficulty and duration.

REFERENCES

Armen A. Alchian and Harold Demsetz, "Production, Information Costs, and Economic Organization", *American Economic Review,* 62, pp. 777 -97, 1972.

[5] The worker's messages still need to specify intervals if the value function has continuous support, but the length of those intervals introduces a loss of surplus in addition to that caused by randomness.

George Baker, Robert Gibbons, and Kevin J. Murphy, “Relational Contracts and the Theory of the Firm”, *Quarterly Journal of Economics,* 117, pp. 39 – 84, 2002.

Ian Ball, Matthew O. Jackson, and Deniz Kattwinkel, “Comment on Jackson and Sonnenschein (2007), “Overcoming Incentive Constraints by Linking Decisions””, *Econometrica,* 90, pp. 03 - 07, 2022.

Ian Ball and Deniz Kattwinkel, “Quota Mechanisms: Finite Sample Optimality and Robustness”, Working Paper, MIT Department of Economics, 2025

Patrick Bolton and Mathias Dewatripont, ”Authority in Organizations: A Survey” pp. 342 -72 in *Handbook of Organizational Economics,* Robert Gibbons and John Roberts (Eds.), Princeton, NJ: Princeton University Press, 2013.

David, Herbert A., and Haikady N. Nagaraja, *Order Statistics,* 3rd Ed., Wiley and Sons: NY, NY, 2004.

Juan F. Escobar and Juuso Toikka, ”Efficiency in Games with Markovian Private Information”, *Econometrica,* 81, pp. 1887-1934, 2013.

Alexander Frankel, “Discounted Quotas”, *Journal of Economic Theory* 166, pp. 396-444, 2016.

Sanford Grossman and Oliver D. Hart, “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration”, *Journal of Political Economy* 94, pp. 691-719, 1986.

Robert H. Guest, “Of Time and the Foreman”, *Personnel*, 32, May 1956, pp 478-86, 1956.

Oliver D. Hart and John J. Moore, “Property Rights and the Nature of the Firm”, *Journal of Political Economy,* 98, pp. 1119-58, 1990.

Matthew O. Jackson and Hugo F. Sonnenschein, “Overcoming Incentive Constraints by Linking Decisions”, *Econometrica* 75, pp. 241-57, 2007.

Jonathan Levin and Steven Tadelis, “Contracting for Government Services: Theory and Evidence from US Cities”, *Journal of Industrial Economics*, 58, pp. 507-41, 2010.

Roy Radner, “Repeated Principal-Agent Games with Discounting”, *Econometrica* 53, pp. 1173-98, 1985.

Jerome Renault, Eilon Solan, and Nicholas Vieille, “”Dynamic Sender-Receiver Games”, *Journal of Economic Theory,* 148, pp. 502-34, 2013

Herbert A. Simon, ”A Formal Theory of the Employment Relationship”, *Econometrica* 19, pp. 293-305, 1951.

Oliver E. Williamson, *Markets and Hierarchies: Analysis and Antitrust Implications,* New York, NY, : Free Press, 1975.