2024

Strategies for the Contracting Game

USING GAME THEORY TO DEVELOP EFFECTIVE PROCUREMENT CONTRACTS BY MARK NADEAUEMKHOS.COM

BY MARK NADEAU

EMKHOS | emkhos.com

Developed by Emkhos, LLC

Abstract:
This paper summarizes some of the key ideas in Game Theory as it relates to contracting. The topics are based on risks in contract planning, specification and Statement of Work (SOW) drafting, defining specificity, setting cost targets, and developing (and managing) inspection and Quality Assurance plans. Each risk is presented with suggestions for practical applications to mitigate the risks. The applications are framed around five principles related to specificity, informativeness, equal compensation, monitoring, and incentives.

Some of this content is based on Mark Nadeau's 2020 publication, Drafting the Procurement Contract: Strategy, Style, Process (ISBN 9798642067284). There are also references to the work of Nobel laureates Oliver Hart and Bengt Holmstrom.

Contracting is a Game
That's right – contracting is a game! Before I try to convince you, let me define some of my terms. The parties to a contract are the Principal and the Agent. The Principal is the one receiving the benefit of the work and is otherwise known as The Owner, The Buyer, etc. And the Agent – or the Contractor - is the one performing the work (and getting paid for that performance). The transactional relationship between the Principal and Agent is known as the Principal-Agent Problem. The problem comes from information asymmetry – neither party ever has complete information. This leads to problems in the period before contract award (the *ex ante*) called adverse selection. And problems after award (*ex post*) are from moral hazard. The Principal-Agent problem and the balance of the competing interests of the parties can be analyzed using game theory.

The rules of this game are based on contract law and the terms of the specific contract that's been developed between the two parties. That is – the law recognizes that two parties are free to establish their agreement in any way that suits them. The courts step in only to enforce the rules that the players have established for themselves. (In that way, the court is a sort of referee at-large for the contracting game.) The best gaming strategies – that is, the strategies that create the best value for a player - respect the Law of Large Numbers, or the long-term statistical averages of events. A player can win a gamble on a single play that has very little chance of winning, but that doesn't mean the player's strategy is successful!

To analyze a Principal-Agent problem using game theory, we respect the fact that each player (party to the contract) can make choices about how to behave. For example, during performance, the contractor is constantly making choices about whether to comply with the requirement to work or to shirk. And the Principal must make decisions about whether to monitor the contractor's performance or just trust that it's being done. There's a cost to monitoring, after all, and it's prohibitively expensive for an Owner to perform 100% inspection of any work. Game theory can tell us the likelihood of these behaviors by modeling the options in a Payoff Matrix.

A simple example might illustrate. Let's say that the Principal contracts with an Agent to guard a door. Keeping guard over the door is important to the Principal because let's say… the Principal is storing a big pile of money behind that door. And there are thieves about! Now, the Principal and the Agent have a contract that requires the Agent to guard the door. But… it's just a contract! The Principal can't be guaranteed that the Agent is always standing guard during the term of the contract. The Principal is worried about the Quality of the Agent's performance and must decide how much monitoring is necessary to assure that the performance is occurring (and to seek remedy for failures).

The Agent, on the other hand, has a stake in this game. Standing guard at the door costs the Agent money. It's expensive, after all, to hire a good guard for an eight-hour shift. Labor laws require the employer to give the guard periodic breaks, so that means hiring another guard to fill in during the shift. Why not just let the guard on duty slip away and leave the door unattended for 15 minutes – what could possibly go wrong? Heck, why not just hire a part-time guard to do six-hour shifts, leaving the door unguarded for two hours every shift? There would still be better than 50% odds that the Owner isn't going to catch the Contractor shirking. And – if you don't catch the moral hazard, it's as though it never happened!

Each party's sense of value in this situation can be represented in a four-square Payoff Matrix like the one below. Each party has a choice – for the Principal it's to Inspect or Ignore; for the Agent it's whether to Shirk or to Work. Inside the matrix is an estimate of the value – the costs and the profits that result from each combination of behaviors.

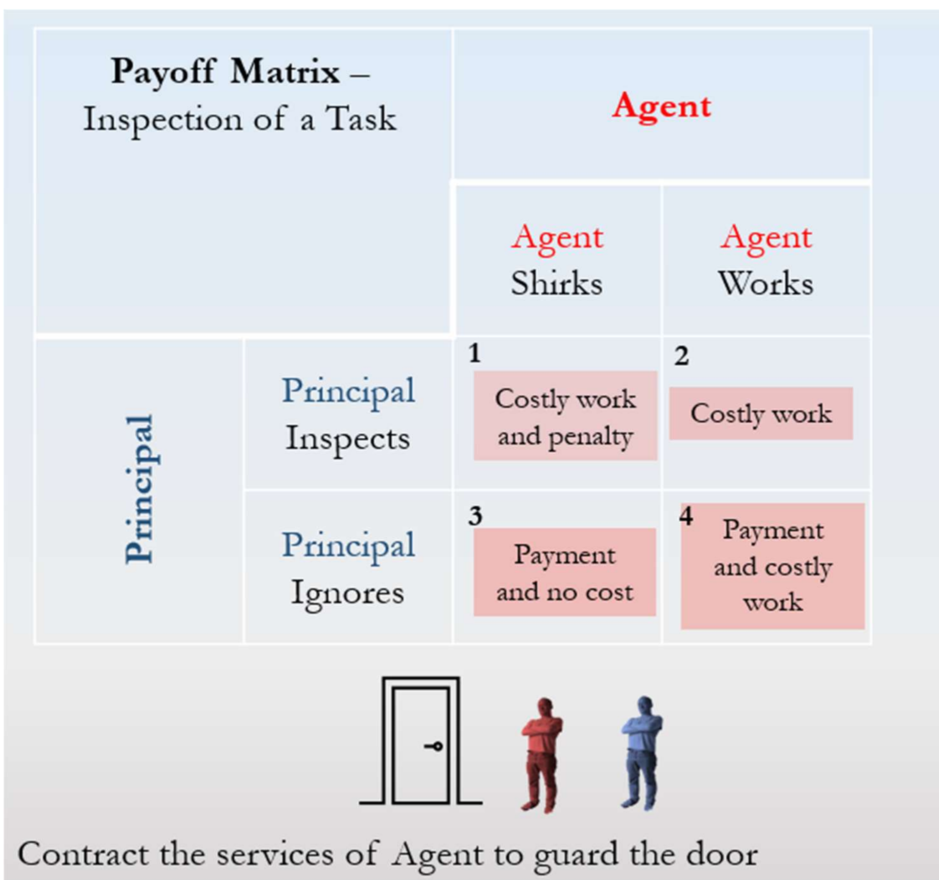

Strategic Design

Using this model, we may be able to design contract mechanisms that steer the Agent to work without needing an inefficient burden of inspection costs by the Principal. In game theory, the Nash Equilibrium is a state in which neither party can benefit from a change in strategy as long as the other player's strategy remains the same. In this case, we want to find what strategy will minimize the Principal's costs. Well, that's simply having the Agent work while the Principal ignores. Such an optimal state can be achieved if the contract establishes the costs and benefits of the other states such that the Agent will maximize profit by working. And – also important – the state will only be achieved if those contract requirements are enforced.

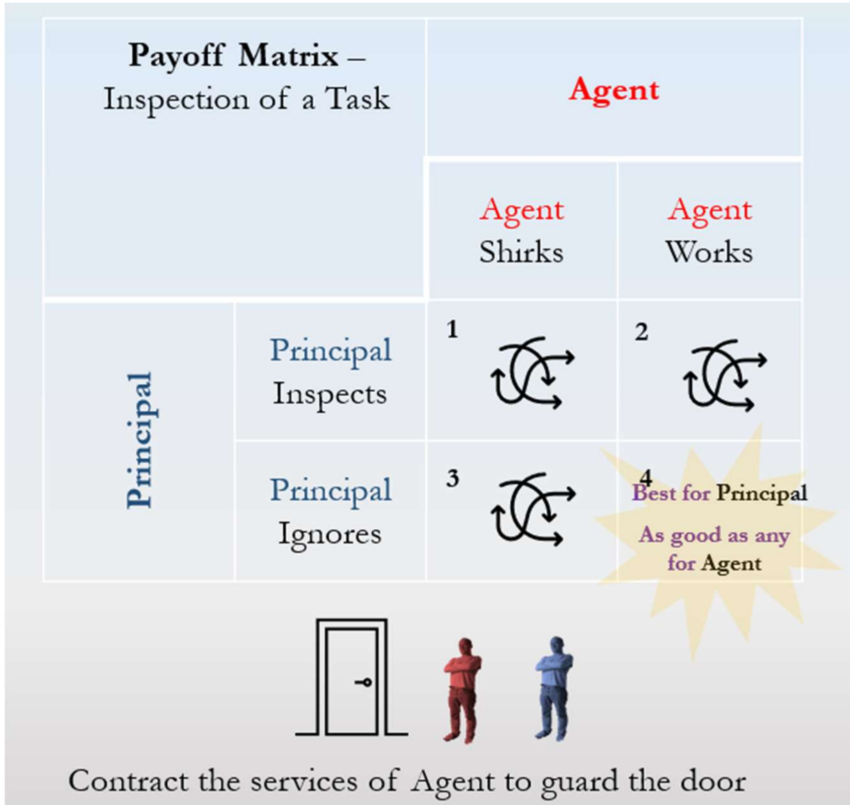

The optimal strategy can be discovered by valuing performance choices in a payoff matrix. Values must be considered carefully, of course. For example, it may be critical to weigh some of the Agent's costs for hosting an inspection against the Principal's value for inspecting, as often there are costs that are quietly passed on to the Principal in the contract. And

Application of game theory to **REVERSE AUCTIONS**

For the bidder: Set a walk away value and stick with it. That is, the bidder's best payoff will come from remaining in the auction until it reaches their valuation and then dropping out. This is like the adage for gamblers to set themselves a limit before they sit at a gaming table and begin gambling.

For the buyer: The Principal should refrain from even using a reverse auction for bids if it's not possible to restrict the number of bidders. That's because the information asymmetry becomes too disadvantageous for the Principal when there are more than three bidders who are somewhat privy to each other's strategies. And the Principal should also set a secret "walk away" value - called Reservation Pricing, establish auction rules for unilateral right to end the auction, and terminate the auction when the price reaches the Reservation point.

players should plan to abide by their valuation throughout the game. The reverse auction, a common type of contract bidding, provides an illustration of the use of these valuation rules.

Cooperation

Solicitation and bidding processes offer a simplified application of game theory. The *ex post* performance of the contract, after award, opens up many other applications. And the iteration of any game introduces complexity. That is, guarding a door is a service over a given performance period. No matter how brief the performance period, the guard doesn't stand at the door for a single moment in time. There are many opportunities to second-guess the strategy, especially when the contractor notices that nobody's watching! The complexity of the contract performance game may be similar to the problem of the hypothetical Prisoner's Dilemma, in which two prisoners make simultaneous decisions whether to rat the other out. It's sort of like flashing your choices in Rock, Paper, Scissors, except the payoff isn't quite as simple as one object cutting, crushing, or covering the other. In the Prisoner's Dilemma, the interrogators are a bit devious, and they offer the greatest reward for a prisoner who chooses to rat while the other chooses not to rat. If they both rat on each other, they both get a meager payout. And if neither rats, they both get a medium payout. Here's a Payoff Matrix showing each prisoner's value of ratting or not. The valuation of the actions are separated by a virgule to show each of the Prisoner's payoff:

Payoff matrix for a prisoner's dilemma		**Prisoner B**	
		Don't rat	Do rat
Prisoner A	Don't rat	3/3	5/0
	Do rat	0/5	1/1

It's clear that to always rat is the safest strategy – it can never be taken advantage of; it's a rock when there's no such play as paper. But brutally ratting out your opponent on every play sacrifices the opportunity for larger payoffs with an opponent who's eager to cooperate. So there is no "best" strategy; it depends on the opponent's strategy, and each player is able to adjust their strategy throughout the game.

Anyway, you probably see that this opportunity for cooperation also applies to the Principal's decision to either Inspect or Ignore and the Agent's decision to Shirk or Work. Following is a summary of the best common strategy using the Axelrod cooperation rules, devised by simulations in which different algorithms compete against each other. That is, here's what makes for the most successful outcome for both parties in a contract relationship:

Application of game theory to **QUALITY ASSURANCE**

1) Be Nice – never be first to defect,

2) Be Forgiving,

3) Be Retaliatory,

4) Be Clear

5) Be cryptic - a known length of game leads to loss of cooperation

Five Principles

Following are some practical rules that will lead to a better strategic position for the buyer (Principal).

The Incentive Intensity Principle
This is about strategic use of incentives to motivate the contractor to a different game state. Payment, which is a fundamental defintion of a contract (work is exchanged for consideration), is one of the handiest tools in this kit. One way to leverage payment is to create a **Shadow of the Future**, which motivates the agent (contractor) to behave in the current state in order to obtain a reward in a later state. This can be structured in an options contract, in which the agent's follow-on work will be offered and awarded at the unilateral whim of the Principal. This option can be written into the contract as an unfunded line item, and Key Performance Indicators (KPIs) can inform the Prinicipal's right to exercise the option.

Cost Targets can also be used as incentives, even though they could be considered disincentives! That is, a contract can be designed such that the agent's fee, profit, and some costs are reduced or eliminated once a certain amount of payment has been made. This is done in Progressive Design Build contracting using a structure called Guaranteed Maximum Price (GMP), which requires the contractor to complete all work beyond the GMP "at cost," or at the contractor's liability. A price structure can also allow a sharing of the cost savings if the work is delivered below the GMP.

Cost Targets

Distribute the profit incentive so that any scope increase is self-defeating to the contractor's opportunism

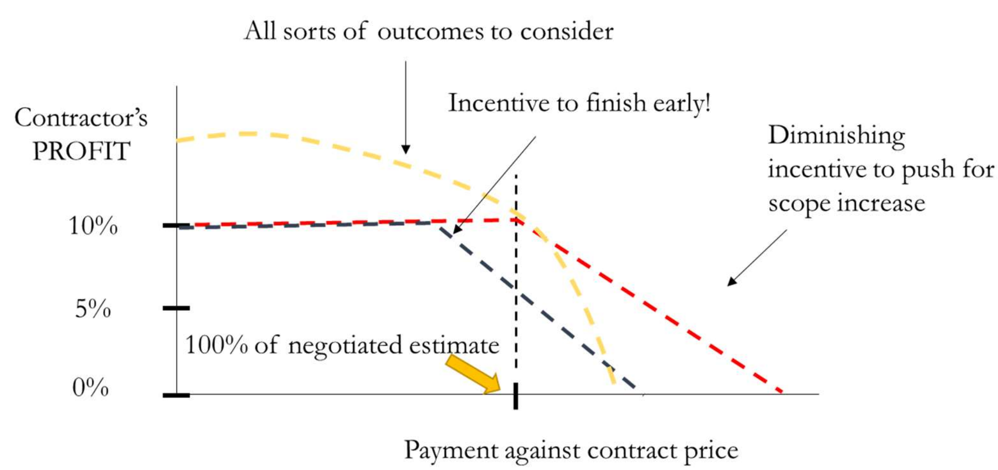

There are monetary incentives in strategic assignment of **Relational Specific Assets**, such as the decision to let the contractor (instead of the Owner) own and manage the IT system that's used in performance of the contract. Both parties should carefully consider their total cost of ownership in this deal, and those considerations should include liabilities for performance of the system and licensing costs.

The Equal Compensation Principal suggests that targeted measures will be favored over those that aren't easily measured or those that are less valued by the Principal. So the buyer should either abstain from favoring any particular task with incentives or accept the fact that performance of some incentivized tasks will suffer. If you have an element that you value but it's very difficult to measure, consider shifting its performance under the umbrella of another type of incentive, such as Relational Specific Assets. (Hmm… I can't figure out how to tell the contractor to ensure proactive maintenance on the building, so let's have the contractor own the building!)

And in strategizing any incentives, keep in mind that they are gameable. If they're not designed carefully, the Principal can end up just giving money away! For example, if you're using KPIs as a criteria for future work or to decide award fees, be sure to establish an honest, unbiased set of criteria and to treat the KPIs to fair evaluation. Another important strategy is to place the incentives where they'll make the most difference. Setting retainage in construction contracts is an illustrative example. If retainage is set as some small percentage of the total price, a contractor might be tempted to walk away from completing the work, even if there's a Liquidated Damage (LD) penalty. The retainage should be set with appreciation for the realistic costs that the Principal will face after Performance Bonds and not necessarily counting the difficult-to-enforce LDs.

The Informativeness Principle basically tells the Principal to share as much information as possible with bidders. It reduces the asymmetry of information among Principal and Agent, and that leads to a more balanced exchange of work for payment. There are some important exceptions, however.

- In a negotiation or reverse auction, do not reveal the reserve or "walk-away" value. And in any negotiation do not reveal the owner's budget, and be careful about revealing any negotiation objectives (balancing how much to say and how much to keep secret is a game in itself).
- In a solicitation, do not reveal to any single potential bidder any information that might give them a competitive advantage. And don't reveal proprietary information until it can be protected through an executed contract.

The Monitoring Intensity Principle brings us back to the Nash Equilibrium problem with which we started this paper. And there are indeed some more specific strategies for the Principal to optimize monitoring costs without risking too much quality failure. A simple prescription for specifying contract requirements is based on the buyer's likely "knowledge" about the quality of the purchased good or service. With a Search item, the buyer knows the quality with near certainty upon receipt. That's the case when you order something out of a reputable vendor catalog and have it shipped to you in a box with a packing slip. You pretty much know what you've gotten before you even open the box! An Experience product is something that doesn't show its quality until after its in operation. A paint job on a structure might be an example – you don't really know how well the surface has been prepped or how thick and even the paint was applied until time and weather do their thing. And – at the far end of the epistemic spectrum is Credence, in which the buyer may never really know the quality of what they've purchased. This is usually true of consultant services. Following is a sort of chart that generally tells you how to specify requirements depending on the likely knowledge of quality:

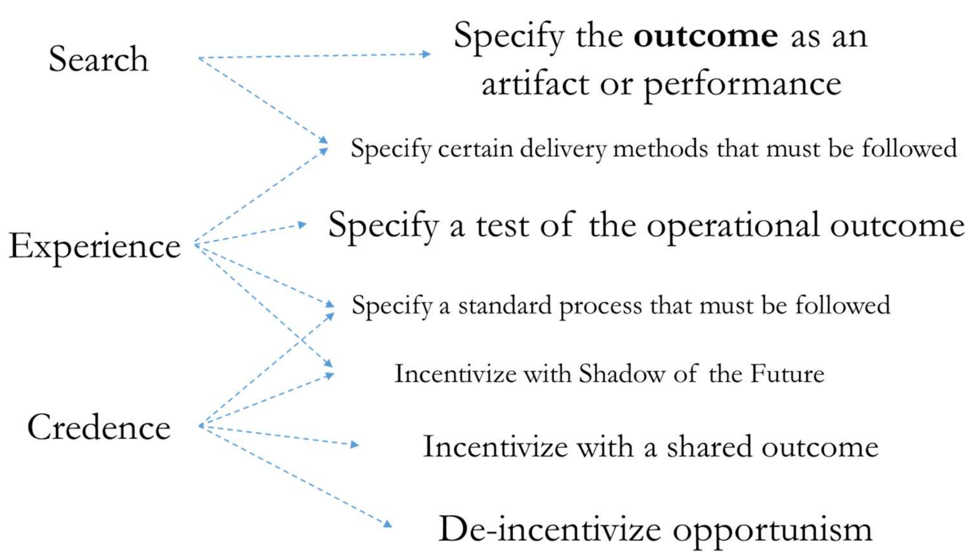

And here's the concept sliced in a different way, based on What we'll know as well as When we'll know it:

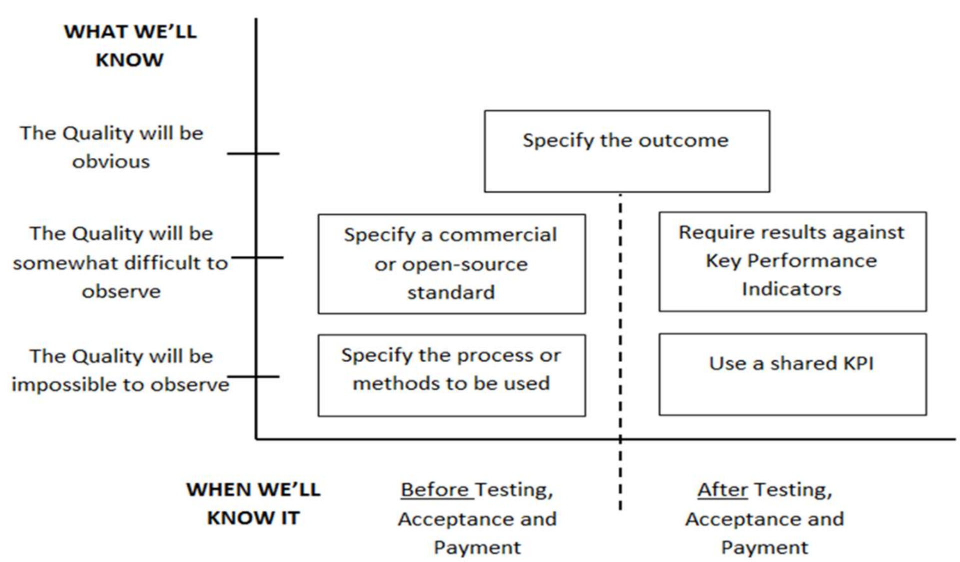

Sample Planning in Quality Assurance

To wrap this up, let's say a few things about how to design a good (cost effective) survey for quality assurance activities. As we discussed with the Monitoring Intensity Principal, the strategy is generally to inspect only as much as you need to; doing too much is needlessly expensive.

The sampling plan should be developed with consideration for the underlying frequency distribution. This is a statistics problem, and anyone designing QA plans should know the basics (and I can't provide that here). With some knowledge of the distribution, it's possible to calculate an ideal sample size. For this sample size equation, it's important to make assumptions about the shape of the distribution and the amount of error that exists in the total population of possibility. A good QA plan incorporates Bayesian logic into this assumption. That is, it's best to update the assumption about the total possible error over time, as the contractor demonstrates their actual performance or as the production run produces consistent results. And Quality requirements should leave some room for error, as it's not reasonable to expect exact conformance. Tolerance for error should be defined as an Acceptable Quality Limit (AQL).

There are some excellent methods for statistically reducing the number of tests that are needed. These include stratified sampling and rank order sampling. I won't go into those here, as I've probably long since worn out my welcome with readers. But hey – if you've made it this far and you're interested in seeing another paper that covers QA sampling plans with some gentle math, please reach out to me. There should be an email given somewhere on the Emkhos.com website. Otherwise, stay tuned to the Emkhos library – I intend to someday add such a paper whether anyone wants it or not!

(the end)