What is the impact of GPOs on healthcare-product supply chains? Recently, researchers at Purdue University used "game theory" to answer this question.[1] Among the specific questions asked and answered:

Do providers experience lower total purchasing costs with a GPO in the supply chain?

The Purdue study concludes that, overall, GPOs decrease providers' total purchasing costs; and, in particular, decrease the total purchasing cost for small providers.

Do contract administration fees (CAFs) mean higher total purchasing costs for providers?

The Purdue study finds that CAFs have no effect on the total purchasing costs of any provider, large or small.

Before explaining and interpreting these answers, the study co-author provides below a little background about game theory. Then co-author reviews all of the major results of the Purdue study.

Finally, in the closing section, HIGPA interviews Professor Leroy B. Schwarz, the study's senior co-author.

A Primer on Game Theory

Game theory is a widely-used technique in economics. The most well-known game theorist is Nobel Prize Laureate John Nash (whose life is the subject of the movie "A Beautiful Mind").

Game theory is used to examine situations in which "players" have competing interests. For example, in the Purdue study, healthcare manufacturers want to maximize profits while healthcare providers want to minimize their total purchasing costs.

As a scientific technique, game theory is like a laboratory test designed to examine the effects of a medication. That is, the conditions in game theory are tightly controlled — even, "unrealistic" — compared to what happens in the real world.

For example, the Purdue study examines a scenario in which providers are seeking the lowest total purchase cost for a *single* product. And, instead of considering a world in which several GPOs compete for providers' business, the Purdue study assumes only a *single* GPO.

It is important to keep these "laboratory conditions" in mind when interpreting the results of any game-theory study. In other words, ask if a result is a consequence of a clearly unrealistic assumption or if something else, more realistic and powerful, is

---

[1] The title of the Purdue University study is "The Impact of Group Purchasing Organizations on Healthcare-Product Supply Chains", currently under peer review at *Manufacturing & Service Operations Management*. The authors are Professors Qiaohai Hu, Leroy B. Schwarz, and Nelson Uhan. Copies may be downloaded here.

causing the results to be what they are. Indeed, by keeping the possibly unrealistic assumptions in mind it is possible to make hypotheses about how the results might carry over to a more realistic scenario. See Curtis Rooney's interview of Professor Schwarz below.

The Purdue Study: The Game

The Purdue study examines a supply chain that has several providers (e.g., hospitals) with different purchasing requirements for a single product; a single manufacturer whose goal is to set its price schedule to maximize its profit; a competitive source selling the product at a fixed price; and a single GPO.

Each provider is free to purchase some or all of its purchasing requirement on-contract through the GPO, direct from the manufacturer, or from the competitive source. In making its decision, the provider's goal is to minimize its total purchasing cost; that is, the cost of the product plus the provider's own contracting cost.[2]

The Purdue study assumes that the provider's contracting cost is the same if it contracts directly with the manufacturer or with the competitive source, but is less if it contracts through the GPO. This savings in contracting cost if the provider buys through the GPO is called the GPO's "contracting efficiency". The GPO charges a contract administration fee (CAF) to the manufacturer on all on-contract sales and a membership fee to providers who chose to use its services.

Note that in the game, it is assumed that the manufacturer's goal is to maximize its profit while the providers goal is to minimize their cost. The Purdue study further assumes that the GPO is a profit maximizer; in particular, that the GPO might charge members a higher price than offered by the manufacturer.

Since it is well known that some GPOs operate on a not-for-profit basis and, more generally, that such "margin-gain" doesn't occur at any GPO, this assumption deserves an explanation.

For years, GPO critics of have accused GPOs of profit-making, regardless of what GPO goals might actually be. So, in constructing the game, the Purdue researchers assumed what the critics have alleged: that GPO's goals are to maximize its profit. This was in order to see how manufacturers and providers might be affected in this "worst-case" scenario.

The Rules of the Game

---

[2] Schneller and Smeltzer (2006) identify the following components of a provider's fixed contracting cost: determining product use and requirements, obtaining user input, accessing supplier lists, preparing bids or requests for proposals, sending requests for proposals, responding to supplier questions, analyzing bid proposals, conducting product evaluation, selecting suppliers, signing contracts, monitoring contract compliance, and monitoring market competitiveness.

Based on case studies supported by Novation, Schneller and Smeltzer report that a provider's fixed contracting cost is $3,116 per contract if a provider contracts directly with a manufacturer (or competitive source), and $1,749 per contract if a provider contracts through a GPO.

Every game has rules. In the Purdue game, there are two important rules. The first involves "perfect information". This means that every player knows *everything* that *every other* player knows. So, for example, the manufacturer knows the goals of the providers and the GPO, the contracting cost of the providers, and the price being offered by the competitive source. In reality, of course, only some of these things are known to some players, but not to others.

The second important rule involves the sequence in which the players make their choices. In the Purdue game, before the game starts, the GPO has chosen the manufacturer and they have agreed on the contract administration fee (CAF). The manufacturer "moves" first by choosing its price schedule. Next, the GPO chooses its on-contract price. Finally, the providers independently decide on whether to purchase direct from the manufacturer, through the GPO, from the competitive source, or some combination thereof. In the language of game theory, the manufacturer is the "leader" of the game; i.e., it initiates the game by announcing a discount schedule. The GPO chooses next; and, finally, the providers decide how much to buy from each source. Everything else being equal, this "tilts the playing field" in favor of the manufacturer.

If this were, in fact, a game, then the manufacturer would get to "go" again; followed by the GPO, etc.: Like a simple around-the-table negotiation, each player taking its turn until the end. In fact, since the manufacturer knows everything that every other player knows, the manufacturer is in a position to anticipate how the GPO and then the providers will "move" in response to *every* possible decision it might make. So, the manufacturer will make a single decision — the one that is best for it — and the GPO and providers will do the same. After this first round, the manufacturer *could*, in theory, "go" again, but, since it made the best possible decision for itself in the first round, if it were allowed to "go" again, the manufacturer would make the very same decision. This is called "equilibrium"

How Does the Game End?

In game theory, as in real life, things either reach an "equilibrium", where every player is satisfied with its last choice; or things never reach an equilibrium. In real life, unless an equilibrium is reached, someone eventually gets tired or disgusted, walks away, and the process simply stops. Game theory is focused on those scenarios when an equilibrium is reached; and then in interpreting what the "equilibrium results" are.

The Findings of the Purdue Study

The Purdue study reports the following equilibrium results:

> **Do providers experience lower prices or lower total purchasing costs with a GPO in the supply chain?**
> *"...the GPO will set its price to be equal to... the price that equalizes the total purchasing cost of the largest provider that (the GPO wants) to buy on contract .... Providers with smaller purchasing requirements will experience lower total purchasing costs in the presence of a GPO, but may experience higher per-unit prices."*

So, as described in the introduction, the Purdue study concludes that, overall, GPOs decrease providers' total purchasing costs; and, in particular, decrease the total purchasing cost for small providers.

Question: But what about the "may experience" higher per-unit prices? And, does your study conclude that large providers don't experience any savings in total purchasing cost?

Professor Schwarz: Remember, the providers are motivated to minimize their *total* purchasing cost. This is like a consumer choosing to purchase an item either at a supermarket or at a Sam's Club. Assume that Sam's Club has a lower price for the item. Does the smart buyer buy at Sam's Club? If the Sam's Club is next door and the supermarket is across town, then, yes! But, what if Sam's Club is across town and the supermarket is next door? Whom to buy from *should* depend, in the language of the Purdue study, on the consumer's "total purchasing cost", which includes the consumer's "contracting cost": in this example, the cost of gasoline, parking, and, of course, the consumer's time, *plus* the cost of the item to be purchased.

Now, suppose that the owner of the local supermarket knows that the consumer's total purchasing cost to buy at Sam's Club is, say, $10: $4 for the item and $6 in contracting cost. In order to attract the consumer to buy from his supermarket, all he has to do is to set his price so that the consumer's total purchasing cost is equal to $10 or slightly less, say $9.99. So, if the consumer's contracting cost to buy at the supermarket is $5 (e.g., less time and gasoline), then the supermarket will set it price to $4.99. The result? The consumer will choose to buy at the supermarket, and incur a total purchasing cost of $9.99, one penny less than Sam's Club $10.00. However, the consumer may pay more per unit ($4.99) than he/she would at Sam's Club ($4).

Roughly the same thing is going on in our study, which has a single GPO (supermarket). Incidentally, this result is consistent with one of the findings of the 2002 GAO study. Namely, that buying through a GPO does not guarantee the lowest unit price.

Question: Is your result about possibly higher unit prices a consequence of one of the simplifications of the model?

Professor Schwarz: Yes, it *probably* is. To see why, let's return to the consumer's decision. Suppose, now, that there are several local supermarkets whose "contracting costs" for the consumer are all about the same. This is like adding GPOs to the Purdue game. Now there will be competition among the supermarkets, each of them wanting the business of the consumer and each willing to reduce its per-unit price to do so. The result: Lower per unit prices for the consumer. So, I believe that adding GPOs to the Purdue game will yield lower per-unit prices for the providers.

Question: What about large providers? The quoted result is that one or more of them won't experience any savings in total purchasing cost in buying through the GPO.

Professor Schwarz: Another factor that the Purdue game ignores is that large providers are often part owners of the GPO, and, hence, will receive share-backs and distributions. Under these circumstances, large providers buying through the GPO will experience lower (net) per-unit costs and, hence, lower total purchasing cost.

Do CAFs mean higher prices paid by providers?

"In the two special cases examined, the total purchasing cost of the providers is not affected by the CAF ... Computational experiments indicate that this behavior occurs in the general3 case as well."

Hence, as described in the introduction, the Purdue study concludes that CAFs have no effect on the total purchasing costs of any provider, large or small.

Question: Manufacturers claim that CAFs increase their cost of doing business and "force" them to increase prices to all providers, whether they buy on contract or not

Professor Schwarz: As one of the members of the study team, I *initially* believed the manufacturers' claim. I don't anymore, based on our results.

CAFs are similar to sales commissions paid by sellers to businesses like Groupon and LivingSocial. Although each operates in a slightly different manner, both businesses organize a group of volunteer buyers to purchase some given product or service from a specific seller. The seller, say a home-furnishing store, agrees to lower its price to each of the buyers in exchange for getting the business of the group. And, the seller pays a 30-50% commission to Groupon or LivingSocial for setting up the deal.

Notice that no one is *forcing* any given seller to sell to the group or to pay a commission to the group purchasing organization. So, returning to your question: I wouldn't be surprised if that home-furnishing store complained that it has to charge its all its customers more because of the commission it "has to" pay to the group purchasing organization. Would you?

Question: So, what does your study suggest would happen if the "safe harbor" provision of the Social Security Law were removed, and GPOs were not allowed to charge CAFs to manufacturers?

Professor Schwarz: Our study doesn't directly address that question. We do address a scenario in which there is *no* GPO, and in it providers pay higher total purchasing costs.

---

[3] The "special cases" assume either only two providers with different purchasing requirements or n>2 providers with identical purchasing requirements, and in either case, that the manufacturer's quantity-discount schedule is linear. For these special cases, game theory yields unambiguous results for all parameterizations. The "general" case involves a nonlinear discount schedule and/or n>2 providers with different purchasing requirement. In the general case, the equilibrium in the game cannot be derived but must be computed.

What affects GPO profits?
*"In the special cases examined, we have demonstrated that GPO profit is
nondecreasing in its CAF and nondecreasing in the GPO's contracting efficiency.
Indeed, for low values of these parameters the GPO makes no profit.
Computational tests display the same behavior for more general cases."*

First, this result says that the higher the CAF the GPO is able to charge the
manufacturer, then the higher the "profit" of the GPO. For not-for-profit GPOs this
means either lower prices for members or higher distributed profits for its member-
owners. It also says that the more efficient the GPO is at contracting for its members, the
higher its "profit". This is an incentive for either for-profit or not-for-profit GPOs to "get
better" at contracting.

Question: Our member GPOs have been saying for years that they are better at
contracting than any single provider could possibly be. Are you saving that the more
efficient GPOs become at contracting, the better off its provider-members will be?

Professor Schwarz: In a word, yes. Our results tell us that, everything else being equal,
a GPO's "profit" will increase the more efficiently it is able to contract for its members.
Higher "profit" means that the GPO is in a position to offer lower prices to all its provider
members or to share its "profits" with member-owners.

**How are supply-chain profits divided between the manufacturer and the
GPO and how is this influenced by the "power" of the GPO?**
*"As displayed in all cases examined, the GPO's share of supply-chain profits are
nondecreasing in both its CAF and its contracting efficiency. The more powerful
the GPO is in negotiating its CAF, and the more efficient it is, the higher its profit
and its share of total supply-chain profit."*

This result says that if one fixes the *total* profit that is made by either the
manufacturer or the GPO, then the higher the CAF and the more efficient the GPO is at
contracting, the more of that total profit goes to the GPO. Naturally, this motivates the
GPO to charge higher, not lower CAFs. Also, as above, it motivates GPOs to become
more and more efficient at contracting for its members.
This result and the result that CAFs don't affect providers' total purchasing cost
are worthwhile considering together. In brief, CAFs don't affect providers' total
purchasing cost, but that CAFs reduce manufacturer profit; and that the higher the CAF,
the lower manufacturer's profits. Among other things, this result explains why
manufacturers would like to reduce or eliminate CAFs.