ARTICLES

ESTIMATING DAMAGES TO DIRECT AND INDIRECT PURCHASERS IN PRICE-FIXING ACTIONS

James A. Brander and Thomas W. Ross*

This paper considers the measurement of damages, due to price-fixing, to both direct and indirect purchasers, focusing in particular on the importance of pass-through – the extent to which price overcharges at the level of producer cartels are passed through to final consumers. We critically review theoretical principles and applied techniques for measuring damages, emphasizing the leading techniques: reduced-form regression analysis and the use of comparator benchmarks. The most common approaches to the assessment of pass-through and measurement of the consequent harm to indirect purchasers are also considered. Several important considerations associated with the use of regression techniques are explored, including issues related to the interpretation of diagnostic statistics, issues of model specification, and the potential for specification search or data-mining bias.

Cet exposé porte sur l'évaluation des dommages causés par la fixation de prix aux acheteurs directs et indirects, mettant l'accent sur l'importance du transfert – soit la mesure dans laquelle les majorations de prix au niveau des cartels de producteurs sont transférées aux consommateurs finaux. Nous examinons d'un œil critique les principes théoriques et les techniques appliquées d'évaluation des dommages, en insistant sur les principales techniques : l'analyse de régression de forme réduite et l'usage de points de référence comparatifs. Les méthodes les plus courantes d'évaluation du transfert et de l'évaluation du préjudice subséquent aux acheteurs indirects sont aussi passées en revue. Plusieurs considérations importantes liées à l'usage des techniques de régression sont analysées, y compris les questions relatives à l'interprétation des statistiques diagnostiques, les questions de spécification de modèles et le potentiel de recherche de spécifications ou de biais de forage de données.

I. Introduction

Dealing with price-fixing and related anti-competitive practices is a central activity for competition policy authorities such as Canada's Competition Bureau. However, private enforcement of competition law (i.e. private lawsuits), particularly through class actions, has become very significant in both Canada and United

States, and is taking on increased significance in Europe.[1] In price-fixing class actions, plaintiffs normally make a claim for damages based on the economic harm done to buyers of the products at issue. Therefore, estimating these damages is a central part of such cases. And, even in cases pursued by competition policy authorities, damage estimation is often undertaken, for example to provide guidance with respect to the appropriate level of fines.

One important aspect of price-fixing class actions concerns whether only direct purchasers can bring an action or whether indirect purchasers also have legal standing. In Canada, this question was settled by a trilogy of 2013 Supreme Court of Canada decisions dealing with class actions, referred to as *Sun-Rype, Microsoft*, and *Infineon*.[2]

In all three of these cases, indirect purchasers were included in the proposed class of injured parties when plaintiffs sought legal certification of the class. For example, in *Infineon*, the defendants consisted of producers of dynamic random access memory (DRAM), most of whom had been fined for price-fixing in the United States and/or Europe. Defendants sold DRAM to buyers including computer manufacturers such as Apple, Dell, and IBM for use in computers and other electronic devices. These *direct purchasers* sold computers and other devices incorporating DRAM to their customers, who are therefore *indirect purchasers* of DRAM. The logic of the indirect purchaser case is that when defendants fixed prices and overcharged for DRAM used in electronic devices, those overcharges were at least partially passed on in the form of higher prices for those electronic devices, resulting in harm to final consumers.

Under federal law in the United States, indirect purchasers are not able to bring class action cases for price-fixing and related anti-competitive practices, but many states do allow indirect purchaser actions.[3] In Canada, conflicting decisions had arisen at the provincial appeal court level, creating uncertainty over the status of indirect purchasers until the Supreme Court decided in 2013 that indirect purchasers could make damage claims.

In Brander and Ross (2006)[4] we provided an overview of some aspects of damage estimation for price-fixing cases. In this paper we review recent developments in damage estimation focusing particularly but not exclusively on indirect purchasers. We pay significant attention to the estimation of *pass-through* – the extent to which prices charged by defendants to direct purchasers are passed through to indirect purchasers.

We also focus on some econometric issues that we believe are important for the estimation of damages, for both direct and indirect purchasers, but that have received relatively little attention in this context.

In Section II we derive a framework for damage estimation, showing explicitly how damage depends on overcharges, pass-through, and at-issue revenues. Section III reviews the major methods for estimating the overcharge and Section IV discusses estimation of pass-through. Section V identifies a method of dealing with pass-through that incorporates estimation of the combined or net effect of the direct overcharge and pass-through into a single step. One important statistical tool used in estimating overcharges and in estimating pass-through is regression analysis. Section VI provides an overview of some principles of regression analysis that are important in damage estimation, focusing on issues that often cause regression analysis to be confusing or possibly even misleading. Section VII contains concluding remarks.

II. A Framework for Damage Estimation

In Canadian law, price-fixing and related anticompetitive actions are based on Section 45 of the *Competition Act*, which states that "Every person commits an offence who, with a competitor of that person with respect to a product, conspires, agrees or arranges (a) to fix, maintain, increase or control the price for the supply of the product; (b) to allocate sales, territories, customers or markets for the production or supply of the product; or (c) to fix, maintain, control, prevent, lessen or eliminate the production or supply of the product." We often describe price-fixing and related anticompetitive actions as arising from a "conspiracy" but, as the Act indicates, any agreement or arrangement that violates Section 45 might be the basis for legal action. However, for convenience, we will often refer to a price-fixing agreement or arrangement as a conspiracy or cartel.

The *but-for* approach is the foundation for most analysis seeking to estimate damages arising from price-fixing or other anticompetitive actions. Using this approach, damages are normally taken to be the difference between what the injured parties actually paid for the products at issue and what they would have paid in the absence of ("but for") the anticompetitive action. And we often use the term "but-for" as an adjective to identify the value that some variable would have taken on in the absence of anticompetitive actions, such as the but-for price.

What the injured parties actually paid is a matter of fact. However, the but-for situation does not occur in actual fact. Therefore, the but-for outcome is sometimes called the "counterfactual" outcome. In price-fixing cases the main difference between the actual outcome and the but-for or counterfactual outcome is usually that prices are, as a result of price fixing, higher in the actual situation – higher than they would have been but for price-fixing.

A. Direct Purchasers

We first consider the case in which all final consumers are direct purchasers, so there are no indirect purchasers. In a direct purchaser case, the price paid by the purchasers is the same as the price received by the cartel members (apart from any sales taxes).[5] The overcharge is defined as the difference between the actual price and the but-for price. The primary measure of damage is taken to be the overcharge multiplied by the actual quantity purchased. In mathematical terms, if we let the actual price of a product be p^A and the but-for price be p^B, then the overcharge per unit is $p^A - p^B$. If the actual quantity sold to injured parties is q^A then the damage, D, is

$$D = (p^A - p^B)q^A. \tag{1}$$

As has been described in many sources, equation (1) understates the actual harm to purchasers as it ignores the loss due to "quantity effects" that arise when, due to higher prices, purchasers buy less than they otherwise would.[6] Any benefit they would have received from that additional consumption is lost. It is, however, much more difficult to estimate this additional loss than to estimate the loss shown by equation (1). Furthermore, some potential consumers who would have bought the product at the but-for price might not buy anything at all at the actual price and therefore are not identifiable. As a practical matter, equation (1) is normally the basis of damage estimation and that is what we focus on here.[7]

An economist using equation (1) to measure damage has three variables to estimate: the actual price, the but-for price, and the actual quantity. The but-for price is the most challenging of these three variables to estimate given that the but-for outcome cannot be directly observed. And even measuring the actual price and the actual quantity may be difficult. It is sometimes convenient to reorganize equation (1) in the following way. If we multiply (1) through by p^A/p^A (which equals 1), the equation can be rewritten as $D = ((p^A - p^B)/p^A)(p^A q^A)$, which simplifies to

$$D = vR \qquad (2)$$

where v is the proportional overcharge, $(p^A - p^B)/p^A$, and $R = p^A q^A$ is the revenue received by the cartel for sales of the product at issue (the "at-issue revenues"). We have defined the proportional overcharge relative to the actual price. It shows the share of the actual price that is due to the overcharge and it must lie between 0 and 1.[8]

Frequently, the same product will sell for different prices to different consumers, even within the same time period, possibly due to quantity discounts, pre-existing contracts or other factors. In addition, there are often multiple variants of at-issue products commanding different prices. For example, DRAM, the at-issue product in *Infineon*, was sold in a wide variety of forms. One advantage of equation (2) is that if there are multiple at-issue product variants and/or if there are multiple price categories for a single product variant, then equation (2) applies to the entire group of products. This claim is relatively easy to see if each product has the same proportional overcharge. However, it is also true even if the different at-issue product variants and/or price categories have different proportional price overcharges provided that the proportional overcharge, v, is taken to be the weighted average overcharge (weighted by revenue) over the various different product categories.[9]

As equation (2) can be applied to a full price schedule for a variety of product types, it follows that neither different product varieties nor different price categories create difficult conceptual or practical problems, provided the necessary data can be obtained. For example, suppose that some buyers pay standard prices and some pay discount prices. If both standard buyers and discount buyers would have paid, for example, 25% less in the but-for world than they actually paid, then the overall proportional overcharge would be 25% and the fact that some buyers pay different prices from others poses no particular difficulty in estimating aggregate damages. Furthermore, even if standard buyers paid a different proportional overcharge than discount buyers, we can still apply equation (2). The damage is often calculated on a period-by-period basis, such as a year-by-year basis, with different prices and different overcharges for different years within the overall class period.

B. Indirect Purchasers

We now consider indirect purchasers. To be as clear as possible, we focus on a case in which there are only direct purchasers and indirect purchaser final consumers. For example, in the *Infineon* case this would mean focusing only on direct purchasers of DRAM such as Dell and final consumers who purchased computers and other devices containing DRAM from Dell and other direct purchasers. A central question in such cases concerns the extent to which overcharges initiated by the defendants are passed through to final consumers. If overcharges are passed through on a one-for-one basis throughout the distribution process, then pass-through is 100% and equation (2) applies to indirect purchaser final consumers. Specifically, indirect purchaser final consumers suffer overcharge damages equal to the overpayments received by the cartel. These overpayments can be measured as the proportional overcharge imposed by the cartel multiplied by revenues received by the cartel.[10]

In its 2013 judgments on price-fixing class actions, the Supreme Court of Canada specified that damages cannot include double counting. Thus we cannot apply equations (1) or (2) to direct purchasers in the class and simultaneously to indirect purchasers: the overall damage claim cannot, for example, include 100% of the initial overcharge at the direct purchaser level and an additional 100% of the same overcharge at the final consumer level, adding up to 200% in total!

To see the importance of considering harm to indirect purchasers, note that in some Canadian class actions, for example *Infineon* again, a large fraction of the Canadian harm might be indirect. It is even theoretically possible that, in an extreme case, there would be no direct harm in Canada at all. These are cases in which most direct purchasers are not resident in Canada and are therefore not part of the class.

To determine the damage to indirect purchasers, estimating the extent of pass-through is therefore an important part of the process. We can express the damage suffered by indirect purchaser final consumers using the following modified version of (2):

$$D = tvR_I \tag{3}$$

where t represents the pass-through rate, v is the proportional overcharge, and R_I is the at-issue revenue received by the cartel that is attributable to these indirect purchasers. Comparing equation (2) with equation (3) shows that the effect of pass-through is to include an additional factor, the pass through rate, in the damage equation. Using equation (3) to estimate damage to final consumers requires three distinct elements – the estimated pass-through rate, the estimated proportional overcharge, and estimated at-issue revenues.

Suppose that direct purchasers in the class pass on 90% of price overcharges to indirect purchaser final consumers. These indirect purchasers would then have a pass-through rate of t = 90% applied to cartel revenues attributable to their purchases. But direct purchasers in the class would also suffer damage as they pass on only 90% of the overcharge and absorb the other 10%.

More generally, using D_D to represent the damage to direct purchasers and D_I to represent the damage to indirect purchasers, the following formula would apply: $D = D_D + D_I = (1 - t)\, vR_D + tvR_I$. In this formula R_D is the amount of revenue received by the cartel arising from payments made by those direct purchasers in the class. Thus vR_D is the total overpayment by the direct purchasers. However, only share $(1-t)$ of this overpayment is actual damage to direct purchasers as the rest of the overpayment is passed on to indirect purchasers. R_I is the revenue received by the cartel that is attributable to purchases made by the indirect purchasers. If all direct purchasers and all indirect purchasers are in the class, then $R_D = R_I$. However, in Canada, R_D would often be much less than R_I as many direct purchasers would be outside Canada and not in the class. If there are no direct purchasers in Canada then the first term drops out entirely and the damage is just tvR_I, as in equation (3) – the pure indirect purchaser case. Similarly, if $t = 1$ (i.e. 100%) then we get back to equation (3) as all damage is passed through to indirect purchasers. It is also possible that R_D will be larger than R_I as direct Canadian purchasers may export their products to indirect consumers in other countries.

In addition to direct purchasers and indirect purchaser final consumers, there may also be indirect purchasers who are not final consumers. In the *Infineon* case, for example, there may be small scale computer assemblers in Canada who purchased DRAM from distributors or other intermediaries and sold assembled computers incorporating DRAM to retailers or to final consumers. We can incorporate class members in this category in much the same way that we deal with direct purchasers and final consumer indirect purchasers, focusing on the share of the overcharge that they pay.[11]

Keeping track of different categories within the class is conceptually straightforward but may be challenging in practice. However, if pass-through is 100% throughout the distribution system, the situation is greatly simplified as we need consider only the damage to final consumers as given by equation (2) , which is $D = vR_I$ in this case.[12]

If equation (2) or (3) is used, it is necessary to estimate at-issue revenue. Typically the starting point is cartel revenue obtained from quarterly or annual financial statements. However, it is likely that financial statements from cartel firms will contain consolidated revenue covering worldwide or North American revenue, although the class action might be restricted to Canada or to one or a few provinces in Canada. Possibly the cartel members will have data on revenue derived from Canada, but province-level revenue is much less likely to be available. Therefore it may be necessary to estimate the revenue for the province on a pro-rated basis. If no other information is available, using relative gross domestic product (GDP) may be a reasonable basis for making provincial allocations. Sometimes, however, a better indicator of relative importance is available. For example, if the product at issue is used mainly in animal feed, then the relative size of the relevant agricultural industries might be a suitable basis for estimating province-by-province at-issue revenue.

We now discuss the other two components in the damage equation.

III. Estimating Overcharges

The proportional overcharge is needed in both direct purchaser cases and in indirect purchaser cases. Much of the economic literature on price-fixing focuses on estimating overcharges. In a series of papers, Professor John M. Connor has carefully reviewed and analyzed studies of price-fixing overcharges.[13]

A variety of methods have been used to estimate overcharges. In Brander and Ross (2006) we provided a list of such methods. A slightly reorganized version of that list is as follows:

a) Older Methods

i. Simple Before and After Studies
ii. Using Marginal or Average Cost as a Proxy for Price

b) Econometric Estimation of Prices

i. Structural Econometric Estimation
ii. Reduced Form Econometric Estimation

c) Comparator Benchmarks – Using Alternative Markets as Benchmarks

A. Older Methods

The simple before and after method is the longest-standing method and its use pre-dates the use of regression and other formal statistical methods in damage estimation. The idea is straightforward. Assuming there is a single product with a single price at any given time, the pre-conspiracy price is taken to be the but-for price. Thus the pre-conspiracy price is compared with the price prevailing during the alleged price-fixing period (the *class period*). If the pre-conspiracy price is, for example, $10 per unit and the cartel price is $20 per unit, then the overcharge is $10 per unit and the proportional overage is 10/20 = 50%. If there is information on a post-conspiracy period, that information can also be used.

In this simple form, the before-after method raises several issues. First, it is often unclear when price-fixing actually begins. A starting date for legal purposes may be determined as the date at which some evidentiary threshold is met, or it may be determined by some technical legal consideration. In either case it is quite possible that the cartel may have operated and raised prices prior to this date. If so, the before-after test will understate the overcharge because the period immediately prior to the class period might also contain overcharges. To be clear, the problem of correctly dating the price-fixing period arises with other approaches as well, but is particularly pronounced here where there is so little other information used.[14]

A second issue to consider is that other factors may cause prices to change. For example, in Canadian cases involving domestic consumers and foreign-based cartels, exchange rate changes may cause price changes. Other relevant factors may include business cycle variables such as GDP or other measures of income and may also include cost-related variables such as the prices of key inputs or technological progress. In other words, it is possible that a simple before-after approach might attribute a price increase to a cartel when it is actually caused mainly by other factors. It is also possible that a before-after test might fail to identify cartel effects.

For example, technological progress may cause sharp declines in cost that would, in the but-for world, cause prices to fall. In a case such as *Infineon*, where the product is DRAM, a cartel might achieve success by keeping prices stable or allowing a slower price decline than would otherwise occur.[15] Thus the cartel might impose overcharges and economic damage on consumers relative to the but-for world, but the before-after approach would not identify any overcharge. The before-after approach is also difficult to apply if new product varieties are introduced after the start of the class period.

For these reasons, simple before-after comparisons are rarely used if methods that address these issues are feasible. In particular, more sophisticated approaches using price information from before, during and possibly after the class period in regression analysis are important and are discussed under the *Structural and Reduced-Form Econometric Estimation* heading below. Even so, simple before-after tests may in some cases be valuable, particularly if more sophisticated methods are not feasible due to data limitations or for other reasons.

Use of marginal cost or average cost (as alternative measures of unit cost) as the but-for price is another long-standing idea, though one rarely applied in practice. In principle, looking at financial statements or other information recorded by cartel members might allow average cost or marginal cost to be estimated. The logic of using marginal cost or average cost is that, under perfect competition with identical firms in the long run, price, average cost and marginal cost will all come to the same level. Therefore, if the appropriate market structure in the but-for world is perfect competition, then using marginal or average cost as the but-for price is reasonable.

However, one important consideration is that the but-for market structure is often not perfect competition. In industries subject to price-fixing, it is more common that the but-for market structure is some form of oligopoly, where price would not necessarily equal average cost or marginal cost (which would normally differ from each other as well). A second major problem is that it is often difficult to accurately measure marginal or average cost. Conceptually, this cost should be the full cost needed to pay for all factors of production, including paying a competitive rate of return to the owners of the firm. Simply looking at out-of-pocket accounting costs will typically understate actual costs.

As with simple before-after tests, using accounting-based estimates of marginal or average cost as the but-for price is now rare.[16] However, cost information can be used as an important component in more sophisticated methods, as described below.

B. Structural and Reduced-Form Econometric Estimation

An important issue in econometrics that comes up in many areas of economics is the distinction between *structural* and *reduced form* estimation. Both of these methods can be used to estimate the but-for price to be used in damage calculations.

Structural estimation starts by specifying an underlying (and usually well-established) theory that explains how some variable of interest is determined. For purposes of damage estimation the variable of interest is normally price. Suppose, for example, that we believe that the but-for market structure would be perfect competition. If so, we could use the supply-demand model, consisting of a supply function and a demand function, to specify the underlying structure of the industry. The supply function shows how quantity supplied depends on price and other factors (such as cost or exchange rates) and the demand function shows how quantity demanded depends on price and other factors (such as income or other business cycle effects).

The supply and demand functions are the structural equations in this example. One approach to structural estimation would be to estimate both these functions. The price and quantity variables are called endogenous variables as they are determined within the system. That is, given any particular values of the other variables, we use the demand and supply functions to determine the values of price and quantity such that the market is in equilibrium, where the quantity supplied equals the quantity demanded. This equilibrium price would be the estimated but-for price. The other factors affecting supply and demand such as cost, exchange rates, and consumer income are called exogenous variables because they are determined outside the model. We treat their values as externally determined data. If we have data on the values of the exogenous variables during the class period (or for various sub-periods, such as years, within the class period), we can then use the model to estimate the but-for prices during the class period.

Structural estimation of prices does not have to be based on a perfectly competitive model of supply and demand. It can also be based on models of imperfect competition, such as the Cournot model or the

Bertrand model.[17] In Brander and Ross (2006) we distinguished between structural models based on perfect competition and structural models based on imperfect competition. Here we put them in the same general category.

In reduced form estimation we do not start by specifying an underlying theory and corresponding structural equations. Instead we simply specify the variable of interest – price in this case – as a function of exogenous variables, such as cost, income, exchange rates, etc. The advantage of reduced form estimation is that it requires less prior knowledge. We do not need to specify a particular theoretical model of market interactions (perfect competition, Cournot oligopoly, etc). Also, less data is required as we do not need to separately estimate multiple structural regression equations.

However, reduced-form estimation also has disadvantages. First, the reduced-form model conveys less understanding about the economic mechanisms at work. In addition, although writing down a reduced form regression equation does not require explicit assumptions about underlying structure, it does of course rely on implicit assumptions that are not clearly specified and that may be poor approximations to reality. In effect, writing down a structural model requires the analyst to be explicit about the underlying economic assumptions being made and imposes consistency requirements on those assumptions. Reduced form estimation does not impose equivalent restrictions.[18]

There is a very large literature in economics regarding whether structural or reduced form estimation is preferred in a given context and we will not attempt to review that literature here.[19] It is, however, important to emphasize that both structural and reduced form modelling are valuable tools and should be viewed a complements rather than as substitutes.

In practice, reduced-form estimation of but-for prices or price overcharges is more likely to be feasible than structural estimation due to less extensive data requirements. In a specific case in which exchange rate issues are not important, a reduced form estimation equation might be as follows:

$$p = a_0 + a_1PF + a_2I + a_3C \tag{4}$$

where p is the price of the good at issue, PF is an indicator (or "dummy") variable that takes on the value 1 during the class period and zero otherwise, *I* stands for some variable related to demand such as household income or GDP, and *C* stands for some variable related to costs (possibly wages or productivity or some cost index).[20] This equation could be estimated using time series data covering time periods before, during, and possibly after the cartel was operating to raise prices. The estimated but-for price for a given time period within the class period would be the value p takes on if *PF* is set to 0 and the other variables take on their actual values during the cartel period. The difference between the actual price and the estimated but-for price would then be a_1. Therefore, a_1 would be an estimate of the overcharge and the proportional overcharge could be easily calculated.

Equation (4) can be viewed as an extension of the traditional before-after analysis in which we address the possible role of other factors by including them in the regression equation. The overcharge parameter a_1 shows us the estimated overcharge after adjusting for changes in other variables.

The problem that we might not know precisely when price-fixing starts still applies and we do need some information regarding when *PF* takes on the value 0 and when it takes on the value 1. However, one possibility is to allow *PF* to take on the value 1 during the price-fixing period, to take on the value 0 for periods that we are confident are not in the price-fixing period, and either drop observations from periods we are not sure about or use intermediate values estimated in some way.

C. Comparator Benchmarks

The fifth category, which we now call *comparator benchmarks*, is another valuable method. This method uses some alternative market or alternative group of firms that is comparable to the market at issue, or the cartel, except that price-fixing is absent. We refer to this alternative as a *benchmark*. This benchmark situation is used to determine the but-for price or price-cost margin that is used to calculate the overcharge. In Brander and Ross (2006) we used the term *analogy methods* to refer to this category. Some U.S. authors have used the term *yardstick* to describe this method although that term is not ideal for non-U.S. jurisdictions that use the metric system.[21]

McCrary and Rubinfeld (2014)[22] distinguish between yardstick and benchmark methods, defining yardstick methods as methods that use alternative markets, and defining benchmark methods as methods that "evaluate prices only in the market at issue, comparing price in the impact period to available prices in the prices before/or after the alleged period of impact…" However, in what is more common usage, Connor (2014) uses the term benchmark method to refer to any method that can be used to estimate a but-for price or overcharge. That is consistent with Brander and Ross (2006) and with our usage here. Thus an alternative or comparator market is one possible benchmark. This comparator market is similar to what in scientific studies would be called a control. Such a comparator should be as similar as possible to the market at issue apart from the presence of price-fixing (or other anticompetitive actions under consideration).

The comparator benchmark market might involve the same product sold in a different geographic market. An example is provided by the British Columbia credit card case, *Watson v. Bank of America Corporation et al.*[23]. The class in this case consists of merchants who claim that both MasterCard and Visa conspired with banks who issue credit cards to raise the fees that merchants pay when they make credit card transactions. The at-issue product in this case is credit card services. At certification, plaintiffs proposed, among other methods, the possibility of using credit card services in other countries as benchmarks.

The comparator benchmark market might also refer to a similar product in the same geographic market. For example, a conspiracy might exist over one industrial chemical but not over other industrial chemicals produced under similar conditions and subject to similar demand conditions. Different chemicals have very different prices in general. However, the proportional price-cost margins could be compared and used to estimate overcharges.

It would also be possible to compare the price trajectories for comparable products. This would be a form of "difference-in-difference" analysis in which we compare the difference between the prices for the cartel's product inside and outside the cartel period with the difference between the prices of the comparator benchmark inside and outside that period. If we observe pre-conspiracy prices and class period prices, the percentage increase in the benchmark prices can be taken as the but-for percentage price increase for the product at issue. If the actual price for the product at issue during the cartel period exceeds the implied but-for price, then the difference is the estimated overcharge. In effect, using the alternative market in this way solves the problem of dealing with other factors that arises when using the before-after method. The underlying rationale is that prices for the product at issue are affected by the same things that affect the prices of benchmark products, including such things as cost changes, exchange rate changes, and business cycle effects.

Even if prices are not available outside the conspiracy period, it may be useful to compare the prices or margins of a defendant firm with those of producers of comparator benchmark products. One interesting use of a comparator benchmark market arises in *Microsoft*, one of the cases considered by the Supreme Court of Canada. In this case Microsoft is accused by plaintiffs of forming agreements with other relevant firms in the supply chain that reduced competition in certain software markets and allowed it to raise prices above but-for levels. This case illustrates several of the points mentioned in this paper previously. First, there are multiple products at issue including Word, Excel, Office, and Windows, and there are many variants of each of these products that were introduced after the beginning of the class period in 1998. Furthermore, different customers may have paid different prices for the same product based on volume discounts, student status, and for various other reasons.

The primary comparator benchmark in this case consists of other software companies producing (primarily) other types of software. As quoted in the Reasons for Judgment of Justice Myers at the certification stage of this case, Dr. Janet Netz, an expert witness for plaintiffs, proposed the following approach (among others):

"I based a second method on a comparison between Microsoft's profit margins to the profit margins of a benchmark group of successful software firms. To obtain Microsoft's prices on the products at issue in the counterfactual world, I calculate the amount by which these prices would have been lower than Microsoft's actual prices in order to generate the profit margin earned by the benchmark firms. The overcharge was then the percentage by which the actual price was above the counterfactual price."[24]

This method, referred to by Dr. Netz as the margin method, uses other software producers as a comparator benchmark and proposes that in the but-for world (i.e. in the absence of Microsoft's anticompetitive actions) Microsoft would have earned the same profit margin, defined as revenue minus cost divided by revenue, as a comparable set of other publicly traded software producers. Using simple algebra it is possible to calculate the proportional price overcharge from the profit margins for Microsoft and for the benchmark firms. As consistent cost and revenue data on

U.S.-based publicly traded companies is available from well-established sources such as Standard & Poor's Compustat database, such calculations are feasible in this case and in many others.

It is also possible to combine reduced-form econometric estimation with information from comparator markets. A potential example is provided by *Steele v. Toyota*[25]. In this case, brought in British Columbia, Toyota was accused of conspiring with its dealers through its "Access" program to fix the price of Toyota automobiles.[26] Without admitting to any fault, Toyota agreed to a settlement in 2015. If the case had gone to trial, reduced-form estimation of price effects and a comparator benchmark market could have been used in combination. The Access program had a specific starting date and a specific end date, so it is possible to identify the alleged conspiracy period. Also, while the program was implemented in B.C. it was never implemented in, for example, Ontario. Therefore Ontario would be a good comparator benchmark.

It would be possible to estimate a regression of the form given by equation (4) for each of the major Toyota models. We could run regressions for British Columbia including the price-fixing indicator variable as the only explanatory variable: $p = a_0 + a_1 PF$. If the coefficient a_1 is positive (and statistically significant) that would indicate that prices during the class period were higher than in other periods. However, as discussed in the section on before-after studies, it is quite possible that this higher price might be due to other factors. For example, exchange rates have a significant impact on car prices, as does the state of the business cycle and the price of related products, such as the price of gasoline.

One possibility would be include these other variables in the regression — exchange rates, business cycle variables, the price of gasoline, etc. But an even better way to control for these other effects is to run the same regressions for Ontario, where the Access program was never introduced. The *PF* variable would take on the value 1 in the Ontario regression just as in the BC regression, even though the Access program did not operate in Ontario. To the extent that higher (or lower) prices during the Access period were due to exchange rate changes, changes in gasoline prices, or changes in the business cycle, that would show up in both regressions. Changes due to the Access program would show up only in the BC regression. Therefore, comparing coefficient a_1 in the two regressions would indicate if the Access program affected prices.

If the coefficient was approximately the same in both regressions, that would suggest that Access had little or no effect.

A slight variation on this method would be to form a variable equal to the price difference between BC and Ontario for a given model and use this variable as the dependent variable in equation (4). As another use of the "difference in difference" approach, we would be estimating whether the difference in price between BC and Ontario was different in the Access period than in other periods. If BC prices exceeded Ontario prices by more during the Access period than in other periods, this would indicate an overcharge.

These examples illustrate three types of comparator benchmarks: other geographic markets, other similar product markets, and other producers operating in the same or closely related markets.[27]

IV. Estimating Pass-Through

The overcharges described in the previous section refer to the amount by which the defendant increases the price above the but-for level for direct sales. For indirect purchaser actions it is then necessary to determine how much of that overcharge is passed through to class members . There are at least two types of firms involved in the pass-through process. Some are pure intermediaries such as distributors, who purchase a product, such as DRAM, from cartel members, and re-sell that product to others. In addition there are firms that use the product at issue, such as DRAM, as one of many inputs in producing another product, such as computers. We refer to both types of firms as "downstream firms" or sometimes as "intermediaries". There are at least four types of evidence that are relevant to pass-through estimation.

A) information on the market structure of downstream firms and the nature of competition in their output markets
B) statements of industry participants
C) transaction data
D) regression analysis of pass-through relationships

The market structure of intermediaries and other downstream firms and how they compete is particularly important in one specific case. If the intermediary industry is perfectly competitive and if intermediaries as a group have "constant costs" so that the supply curve is horizontal, and the demand curve slopes downward, then pass-through must be 100%.[28]

Even if the downstream industries do not meet the strict textbook criteria for perfect competition, it is common for pure intermediaries (such as major retail chains) to be highly competitive, in which case we expect pass-through to be close to 100%.

If intermediaries are not highly competitive and instead have significant market power, then pass-through may still be 100%. However, it is also possible that pass-through could be less than 100% or more than 100%.

In this context, pass-through of more than 100% means that an overcharge of $1 at the cartel level leads to an overcharge exceeding $1 at the final consumer level. For example, if a retailer experiences a $10 increase in the price of some product and raises its retail price by $11, then the pass-through would be 110%.

More generally, pass-through rates can range from zero to rates far exceeding 100%. For example:

(a) As noted above, if the downstream market is perfectly competitive and producers there have roughly constant average or unit costs, the direct purchasers will already be selling at close to those average cost. This means that any price increase imposed on them by the cartel must be fully (100%) passed on or they will actually be suffering losses. Competition will keep them from passing on more than 100% of the original overcharge.

(b) If the downstream market is a profit-maximizing monopoly with constant unit costs and facing a demand curve with a constant elasticity with absolute value ε, the pass-through rate will equal: $\varepsilon/(\varepsilon-1)$; which will be greater than 100%.[29] A similar situation arises when downstream firms use simple mark-up rules of thumb when determining the prices they charge their customers.

(c) If the downstream market is highly competitive for price increases – perhaps because any higher prices would lead to entry of products from other markets – but the existing downstream firms are making profits at current prices, those firms may not be able to pass on any cost increases and the pass-through rate will be 0%.

(d) If the downstream firm is a profit-maximizing monopoly with constant unit costs and facing a linear demand curve, the pass-through rate will be 50%.[30]

In addition to using the market structure and competitor behavior of the intermediary sector to shed light on pass-through, it is often possible to obtain statements from representatives of intermediaries indicating how they handle price changes or from cartel executives indicating their understanding of the extent of pass-through.

Economists seeking to estimate pass-through will typically use regression analysis if suitable data is available. The pass-through question can be viewed as asking how a one dollar increase in the price paid by a downstream firm affects the price charged by that firm to its own customers. For example, we might ask how the retail price of software carried by Best Buy changes when the price paid by Best Buy for software changes by one dollar.

It might be possible to get data on multiple transactions for a single product, and observe variations in the downstream firm's cost of buying that product and corresponding variations in the price charged by the downstream firm. If so, then pass-through can be estimated using a regression equation of the form:

$$p = a + tc \tag{5}$$

where p is the selling price of the downstream product and c is the downstream firm's cost of acquiring that product from a distributor or from a cartel member. The coefficient t is the estimated pass-through coefficient. An estimate of $t = 1$ corresponds to pass-through of 100%. From equation (5) it is clear that if $t = 1$, then a one dollar increase in acquisition cost, c, would cause the retail price p to also rise by one dollar. However, the overall price would exceed the acquisition cost provided that a is positive, as would normally be the case.

If the downstream firm sells many different products there is no need to restrict the analysis to a single product. For example, suppose that a downstream firm sells many different types of software. If, for each product, we have its acquisition cost and its selling price, then the regression will show us how higher costs translate into higher prices for software generally. It is even possible to include products in different product categories, such as computer software products, computer hardware products, and bundled (combined) hardware and software products. However, we would suggest considering use of indicator variables in the regression when using different product categories if possible. The indicator variable for a particular product would take on the value one for transactions involving that product and zero otherwise. In this context these indicator variables are often called *fixed effects*. Including fixed effects helps to control or adjust for variation across products in the fixed component of the price-cost margin.

The above method can work well for intermediaries such as distributors who are essentially just re-selling a product. In such cases, the acquisition cost of the item represents most of the selling price of the item so it is relatively easy to identify the effect of changes in the acquisition cost.

However, some intermediaries use the at-issue product as an input to produce a more complicated final product. The at-issue input might play only a small role in the final product as, for example, with DRAM included in a computer or television set. If so, other factors might make it more difficult to identify the pass-through coefficient. In such a case it might be preferable to use *all* input costs as the explanatory variable for price, not just the input cost of the at-issue product. And it might be helpful to explicitly include other variables that might affect the price of the final product, such as income. Therefore we might estimate an equation like:

$$p = b_0 + b_1 I + tc \qquad (6)$$

in which p represents the prices charged by intermediaries, I captures variables that affect the demand for the downstream firm's product and c is a measure of all unit costs of the downstream firm. We can estimate this equation using data from inside or outside the price-fixing period. The estimate of the coefficient t will then measure the normal relationship between the downstream firm's costs and its prices. It may be reasonable to assume that cost increases due to the price-fixing of inputs will lead to price adjustments in the same way as any other kind of cost increase. We can then determine the effect on downstream prices by multiplying the change in unit costs caused by the price-fixing by the estimated value for t to get the change in downstream price attributable to the price-fixing upstream.[31] An estimated value for t less than one would then indicate less than 100% pass-through.

In practice, it is often the case that reliable individual transaction level data is not available. A useful alternative might involve using monthly averages. Thus the cost variable would be the average cost across all units of a particular product sold in a given month and the price variable would be the average retail price across those units of the product.

V. One-step Estimation of Multi-Level Price Effects

Application of equation (3) (the damage equation) for indirect purchasers requires two important steps: estimation of the overcharge imposed by the cartel for sales to direct purchasers and estimation of the extent to which that overcharge is passed through to indirect purchasers, particularly final consumers. However, it is possible in principle to estimate the effect of cartel overcharges on final consumers in a single step. This single-step estimation incorporates multi-level effects of the overcharge.

This one-step estimation is based on an equation like (4). However, in this case the price to be explained is the final consumer price. The *PF* variable again identifies the dates of the class period while the *I* and *C* variables represent other factors that affect the demand and costs of the downstream product. A statistically significant coefficient on the *PF* variable would indicate that the upstream price fixing had an effect on prices further downstream. This coefficient estimates the net effect of price-fixing on final consumer prices incorporating the effect of pass-through.

This one-step approach is rarely used, however. Two key problems are easy to see. First, if there are multiple levels in the supply chain between the price-fixed product and the consumer product, it is likely that the price-fixed product may represent a small fraction of the total cost of producing the downstream product, as with DRAM in a large computer or television set. We discussed this situation in connection with pass-through but, provided that intermediaries adopt similar pass-through practices for all inputs, this situation does not necessarily create problems in estimating pass-through, as we can use all costs to estimate pass-through. However, it is a major problem for one-step estimation of price effects as it may be very difficult to see in the data how formation of a cartel for one small product, such as DRAM, affects the overall price for a product like a television as there are too many other factors affecting the final price of televisions.[32] A second difficulty with one-step estimation is that the timing of price adjustment decisions of, for example, retailers of products like computers and TV sets might not coincide with the dates of the price-fixing conspiracy. Retailers might adjust prices with a lag and their prices may stay higher for some period of time after the conspiracy has ended. These lags could be due to contractual commitments to their own customers, the need to honour recently advertised prices for some period or even a general desire to revise prices infrequently.[33]

VI. Principles and Pitfalls in Regression Analysis for Damage Estimation

In both pass-through estimation and estimation of overcharges, regression analysis is an important tool. There is a very large literature on regression analysis and other areas of econometrics addressing many issues that are relevant in damage estimation.[34] In addition, there are several useful surveys dealing with the use of econometrics in competition-policy-related litigation support.[35] We do not review that material here, but we address three specific issues that we believe are important in understanding the role of regression analysis in damage estimation.

A. Reporting Regression Results

When regression results are reported in litigation support or in competition policy proceedings it is normal to present an estimated regression equation along with various statistics sometimes called *regression diagnostics*. These statistics may include standard errors, t-statistics, p-values, R-squared statistics, and confidence intervals.[36] We illustrate the use of estimated regression equations and regression diagnostics using a hypothetical example of pass-through estimation based on DRAM.

We consider two hypothetical intermediaries in the distribution chain for DRAM. One downstream firm, Firm 1, is a distributor that purchases DRAM from DRAM producers and resells it to computer assemblers and retail sellers of DRAM. The other downstream firm, Firm 2, is a seller of custom security camera systems. It buys DRAM from DRAM producers and uses it in security cameras that are sold to final consumers. The pass-through question for each of these firms is: How much does their output price change when the price of DRAM changes?

We focus on just one particular DRAM product type, which is the same for both firms. Suppose we have 16 months of data for each firm. For Firm 1 we have data for each month on the average acquisition price and the average selling price of DRAM it sells to its customers that month. For Firm 2 we have monthly data on the average acquisition price for

DRAM and on the average sale price of security cameras sold. (This is the kind of data that may be available under discovery.)

For each firm we use a linear regression, which means that we assume that the selling price of that firm's product is determined in the following way.

$$p = a + tc + e \qquad (7)$$

where p is the selling price of the firm's product, a is a constant, t is the pass-through coefficient, c is the acquisition price of DRAM, and e is random error term whose expected value is 0. For our purposes here, any factors aside from the acquisition cost are assumed to be incorporated in the random error, e.

We generated simulated data using MS Excel for each firm. For Firm 1, which is simply re-selling DRAM, the acquisition cost of the DRAM accounts for most of the selling price. For Firm 2, the cost of DRAM is only a small part of the total cost of the security camera and is therefore much smaller than the price of the camera. The resulting data and estimated regression lines are shown in the following diagram.

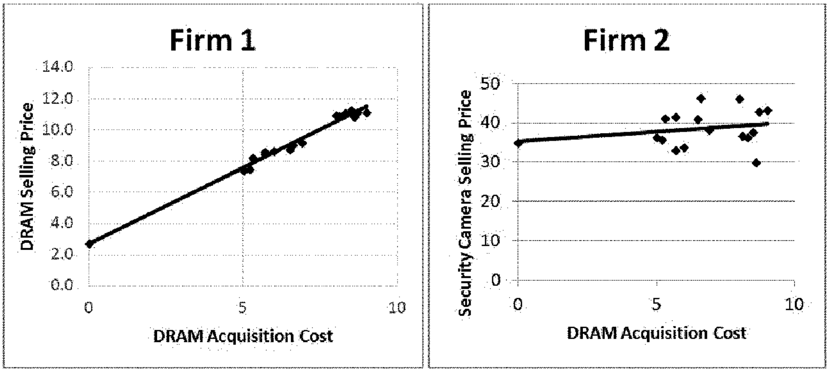

For Firm 1, the estimated regression line is p = \$2.71 + 0.98c. This means that the estimated pass-through coefficient is 98%. The standard error for this coefficient is 0.05, the t-statistic for the coefficient is 18.9, and the p-value is less than 0.001. The reported 95% confidence interval for the pass-through coefficient is 86.6% to 108.7%. The t-statistic, standard error, p-value and 95% confidence interval are all tools to measure the precision with which a coefficient is measured and the likelihood that the true value is greater than zero. Even very small estimated coefficients can be statistically significantly different from zero if they are estimated precisely enough. In this case the small standard error and the fairly tight 95% confidence interval indicate that the pass-through coefficient is estimated with considerable precision. The t-statistic and the p-value indicate that we can be very confident that the pass-through coefficient differs from zero. Similarly, because the 95% confidence interval does not contain the value zero we would say that the pass-through coefficient is significantly different from zero at the 100% – 95% = 5% significance level. The R-squared statistic – a diagnostic that tells us what fraction of the variance in the price the model can explain -- for this regression is 0.96, which is a very high value as R-squared statistics must lie between 0 and 1.

The regression for Firm 1 is a highly informative regression, indicating that the point estimate for pass-through of 98% is likely to provide a good estimate of what is actually happening — of the underlying true situation. The fact that the standard error is small compared to the coefficient estimate, while the t-statistic is relatively large and p-value is relatively small indicates that this is a useful regression, as does the statement that the coefficient is significant at the 5% level. The confidence interval suggests that the true value of the pass-through coefficient is likely not far from the point estimate of 98%. The regression diagnostics are consistent with the figure , which shows that the observed values are very close to the estimated regression line.

For Firm 2, the estimated regression line is p = $35.5 + 0.44c. This means that the estimated pass-through coefficient is 44%. However, the standard error for this coefficient is large, 0.87, the t-statistic for the coefficient is 0.51, which is small, and the p-value is about 0.62, which is large. These values indicate that this coefficient is not significantly different from zero at the standard 5% significance level, and the reported 95% confidence interval for the pass-through coefficient is very wide, covering the range -142% to 231%. The R-squared statistic 0.02, which is very low.

The regression for Firm 2 is very uninformative. It provides a point estimate for the regression coefficient of 0.44 but this is not statistically significantly different from zero, and the confidence interval is so wide as to provide no helpful information at all as it includes both 0% and 100% and a lot more besides. The very low R-squared statistic tells us the regression explains very little of the overall variation in price and that omitted variables (implicitly in the error term) explain almost all of the variation.

An economist would be justified in saying that the regression for Firm 1 is strongly supportive of pass-through being close to 100%. The correct statement about the regression for Firm 2 is that it provides virtually no information about pass-through. However, it is important to understand that the regression for Firm 2 does not provide strong evidence *against* the existence of pass-through. It is uninformative rather than rejecting the presence of pass-through.

As noted above, the data shown here was simulated. In fact it was generated using equation (7) in Excel with particular values for the constant and for the pass-through coefficient, and using the Excel random number generator to generate values for the error term. The main difference is that the error term for Firm 2 was specified to have a much higher variance than the error term for Firm 1. The "true" underlying pass-through coefficient in both regressions is, however, 100% – that is the number we used to generate the data.

The first regression does a very good job in estimating the correct level of pass-through, especially considering that we have only 16 observations. If we had more data, such as 50 or 60 observations, or perhaps hundreds of observations, we would very likely estimate a pass-through coefficient of almost exactly 100%. This is what we would expect in a situation of the type we are trying to represent with this example. Firm 1 is simply reselling DRAM. By far the most important determinant of the price it charges for DRAM is the cost it must pay for DRAM, so the error term (reflecting other variables) would not have much impact on the numbers. The industry is highly competitive so the firm cannot charge much more than the underlying cost and it cannot charge less without going out of business. It would be no surprise that the firm passes through almost exactly 100% of its costs.

For Firm 2, however, DRAM represents only a small part of the overall cost of the camera. This firm produces custom cameras, so every camera is different. Two cameras with the same DRAM might have very different lenses and other different features and therefore have very different prices The error term, reflecting these other influences, is large and important. Most of the difference in price will reflect these other features. To stress an important point – the fact that the regression for Firm 2 did not work well in terms of finding a statistically significant relationship does not mean that the cost increases were not passed on. As indicated, the underlying true relationship that generated the data involves exactly 100% pass-through. A larger sample of relevant data points or a more complete specification of the other contributors to the downstream firm's costs would improve our chances of detecting and measuring that relationship. As previously discussed, an often preferred approach to estimating pass-through for Firm 2 would be to include all costs as the explanatory variable in equation (7) rather than just the cost of DRAM.

The regression for Firm 2 did not find meaningful evidence of pass-through because the error term was large in magnitude for most observations, implying that unobserved factors explain most of the variation in the price of cameras. The regression for Firm 1 did estimate pass-through effectively because the error term was small in magnitude for most observations, as would arise if the main source of variation for Firm 1's DRAM selling price was the acquisition cost of DRAM.

We could also consider cases in which the actual pass-through is small. In such cases we might have good enough data to estimate the low pass-through rate with a high level of precision. For example, we might estimate a pass-through rate of 10% with a 95% confidence interval going from 5% to 15%. Such a case would be a positive finding of low pass-through. That is very different from the situation with Firm 2, where the regression tells us little about pass-through one way or the other.

What happens if the pass-through coefficient is literally zero? How would we distinguish between that case and the case of Firm 2. In both cases we would fail to observe a statistically significant pass-through coefficient. We need to apply some judgement. If we have a very large sample (say 1600 observations instead of just 16) and still fail to find a significant pass-through effect, that would be more suggestive of little or no pass-through, especially if we are able to include in the regression the variables that do explain most of the variation in product price.

This discussion is related to Type I error ("false positives") and Type II error ("false negatives"). In the case of Firm 2, we are making a Type II error. Using standard statistical tests we fail to find statistically significant pass-through even though the "true model" is based on pass-through of 100%. That is a Type II error. A Type I error would arise if there were no pass-through in the "true model" but random variation in the data led us to conclude that statistically significant pass-through is present. Ideally, we would like to reduce the likelihood of both types of error if possible. Having more data and having better data are important in reducing the likelihood of these errors.

If we have limited data so that the likelihood of Type II error is high, then failure to find statistically significant pass-through is not very meaningful. If we have a carefully designed experiment and have a lot of data so that the likelihood of Type II error is low, then failure to find statistically significant pass-through is very significant.

B. Regression Specification

Regression specification refers primarily to the functional form of the regression and to the set of included explanatory variables. In the previous subsection, we have specifications that we know are correct because they are based on the process we used to generate the data. In that example, the main question concerns how accurately we can estimate the pass-through coefficient given the random variability in the data. For Firm 1 we can estimate the coefficient very accurately even with only a small amount of data because the random variability is low but for Firm 2 the regression analysis is essentially useless given the small amount of data because the random variability is high. However, even with firm 2, the specification is correct.

Incorrect specification is another possible source of problems in regression analysis. To illustrate this point we consider another example that often comes up in damage estimation, especially in Canada. The issue is the relationship between Canadian and U.S. prices. Defendant firms in price-fixing cases are often large multinationals that produce for worldwide markets. For example, in *Sun-Rype*, one of the class action cases considered by the Supreme Court in 2013, the defendants included Archer, Daniels, Midland (ADM); Unilever; Cargill and other large producers of various food items, including the at-issue product, high fructose corn syrup (HFCS). Before being litigated in Canada, this case was litigated in the United States.

In *Sun-Rype* and in many other price-fixing cases, significant legal findings and other analysis of prices for the U.S. market is undertaken. To what extent can such information be used in Canada? Plaintiffs may argue that the prices are essentially North American prices with no meaningful difference between Canada and the United States apart from straightforward exchange rate adjustments. Defendants might make the opposite claim – that Canada is a different country with, possibly, very different pricing. Defendants might argue that plaintiffs should start all over in generating price evidence rather than relying in part on U.S. findings. This issue is particularly important in indirect purchaser cases where it is the retail prices that matter and it is perhaps plausible that Canadian retail prices might differ meaningfully from U.S. retail prices in some cases.

Evidence on this issue might include statements from industry participants about the nature of pricing in the industry. In addition, the relationship between U.S. prices and Canadian prices can be investigated using regression analysis. The basic regression question is whether Canadian prices in Canadian dollars can be explained by U.S. prices in U.S. dollars. If the markets are perfectly integrated then the exchange-rate adjusted price would be the same in both markets. This relationship is called the *Law of One Price* (LOP). If this law holds for high-fructose corn syrup (HFCS), for example, then the Canadian dollar price would equal the U.S. dollar price multiplied by the exchange rate, expressed in Canadian dollars per U.S. dollar. For example, if the U.S. price in U.S. dollars is US$100 and the exchange rate is C$1.40 per U.S. dollar then, if LOP holds, the Canadian dollar price would be 100 × 1.40 = C$140. Thus a Canadian indirect purchaser could pay $140 in Canada or could take that C$140, convert it to US$100, and buy the same amount of the product in the United States.

To test whether the Law of One Price holds, the correct regression equation is based on $p^C = xp^U$ where p^C is the Canadian price, x is the exchange rate, and p^U is the U.S. price. It is convenient to take the natural logarithm of both sides, and include a constant term to obtain:

$$\ln(p^C) = a_0 + a_1 \ln(x) + a_2 \ln(p^U). \tag{8}$$

This functional form is sometimes called the *log-linear* form because it is linear in the logarithms of the variables, not in nominal levels of the variables themselves. If the law of one price holds exactly, then a_0 would be zero, and a_1 and a_2 would equal one.[37] However, it would be possible to run a regression using a specification that is linear in the nominal levels of the variables: $p^C = b_0 + b_1x + b_2p^U$. On the surface this appears to regress the Canadian price on the U.S. price while correcting for the exchange rate. However, this specification is incorrect. Therefore, even if the markets were closely integrated such that equation (8) provided a good fit to the data and the regression coefficients were close to the values implied by the law of one price, it is quite possible that the regression in nominal levels would not fit the data well and that the coefficients might not be statistically significant. The problem in this case is that the nominal linear functional form does not reflect the correct relationship between the variables if the Law of One Price holds. An economist employed by plaintiffs might estimate equation (8) and argue that the Canadian and U.S. markets were highly integrated and that the prices in the two countries tended to move together. However, an economist employed by defendants might run the nominal linear regression, find little relationship, and argue that prices in the two countries are not closely related.

The key point here is that it is important to use a suitable functional form for regression analysis. In practice we do not expect to often have specifications that are exactly correct. But they do need to be good approximations. Typically it is not clear what the best functional form is, but regression diagnostics can often be used to help select the best functional form. However, trying different functional forms runs the risk of causing specification search bias, as described in the next subsection.

C. Specification Search Bias

If we do not know the correct functional form it is normal to try several different possibilities and pick the one that fits the data best. It is also normal to try different explanatory variables in the regression and pick the ones that seem to "work best". And other variations in specification can also be tried. For example, in a regression seeking to determine whether the Toyota Access program affected the prices of cars sold, we might try including the price of gasoline, business cycle variables, interest rates, and other variables in the regression. We might also try different ways of organizing the data. We might, for example, try to put all Toyota vehicles in the same regression or we might run a separate regression for each model and emphasize the models that give the "best" results. We might also try using lagged values of certain variables and many other variations.

This process is sometimes called *data-mining* or *data-snooping* or *specification-searching*. However, the term data-mining is also used in computer science to refer to methods for uncovering patterns in large data sets, and the term data-snooping is far from self-explanatory, so we prefer the term specification searching. The advantage of specification searching is that, properly done, it will normally lead to a specification that is a better approximation to reality than the initial proposed specification. However, there are two main disadvantages. First, even when properly done, the process of specification searching can lead us to overstate the confidence we should have in our results. Second, specification searching is prone to misuse – using it to search for results that the analyst wants to get rather than to search for the most accurate specification. Bias in coefficient estimates and significance levels arises when a data set is used for two distinct purposes – first to select a specification, and then to perform estimation and hypothesis tests on the same data set. If this two-step process is used (as is very common and often unavoidable) then some adjustment or reinterpretation may be called for.

These points can be illustrated using the Toyota example. Suppose we have data on Toyota Corolla transactions from before and during the class period (the alleged price-fixing period) for BC . There are many different varieties of Corollas, however, including four-door sedans, coupes, sport models, etc. and there are various options as cars may come with or without air conditioning, with or without a special trim package, etc. We could estimate a regression of the form $p = a_0 + a_1PF + a_2I + a_3C + b_1FE_1 + b_2FE_2 + \ldots$

This regression is very similar to equation (4) except that we have added terms of the form b_iFE_i as fixed effects (indicator variables) that can be used to control for different model varieties and options. This regression could be estimated for the entire set of BC transactions and the coefficient a_1 provides an estimate of the size of any overcharge due to price-fixing. When asking whether this is a statistically significant overcharge we normally use the 5% significance level. If we find an overcharge due to price-fixing at the 5% significance level, this means it is unlikely that we would conclude there was a price-fixing effect if it was not present. It could happen by chance – if we happened to get a lot of prices in the class period that were high for random reasons unrelated to price-fixing. However, the chance of that happening (a Type I error) would be less than 5%. If we find that a_1 is not significantly different from zero, this means that the regression does not support the existence of an overcharge at the 5% significance level.

An alternative procedure would be to estimate a separate regression for each major model variety. Suppose there are 10 major model varieties. In this case we simply estimate equation (4) on ten different subsets of the data, dropping the model variety fixed effects. We would now be quite likely to find an apparently significant price overcharge at the 5% significance level for at least one model variety even if no price-fixing effect exists, just by chance. In fact, if we keep trying different subsets of the data we are virtually certain to find an apparently significant effect sooner or later.

The situation is like tossing a coin. If, for example, we toss a fair coin 5 times, it is unlikely that we would toss heads 5 times in a row. In fact, the probability is less than 5%. So if we do pick up the coin and proceed to toss heads five times in a row we can reject the hypothesis that the coin is a fair coin at the 5% significance level. However, even if the coin is a fair coin, if we keep the tossing the coin, sooner or later we are virtually certain to toss heads five times in a row. We cannot focus on just those five tosses and claim the coin is biased. If we do 10 different trials, tossing the coin 5 times in each trial, we are actually quite likely to get 5 heads in a row in at least one of those trials.

It would be an error to conclude that the coin is biased just because one of those trials generated five heads in a row. Similarly, it would an error to conclude that the Access program led to higher prices just because one out of several models exhibited an apparently significant effect at the 5% significance level. Suppose the model with this effect was Toyota Corolla 4-door sedans with air conditioning. We would not even be justified in concluding that this model variety was subject to overcharges at the 5% significance level just because this regression, taken in isolation, generated an apparently significant effect at the 5% level. The problem is that if we try enough different versions of the regressions we are very likely to find one with apparently significant effects just by chance, just as we are likely to toss five heads in a row if we do enough trials.

In the Toyota case, suppose that the one model variety that generates an apparently significant result yields a 95% confidence interval for a_1 that goes from \$150 to \$250. A statement that the 95% confidence interval goes from \$150 to \$250 can be easily misinterpreted. It may sound as though the probability that the true value of the overcharge is between \$150 and \$250 is 95%. However, that is incorrect. If we are using multiple tests (i.e. with many different models) and picking the "most significant" one, the true 95% confidence interval is much broader. Much of the regression output, such as p-values and confidence intervals, is conditional on having the correct specification in advance, not on selecting the specification (or the model variety we want to look at) on the basis of a first stage process. But we rarely know the correct specification in advance and must try different possibilities.

The more variations that we try, the more likely it is that we will observe apparently significant effects by chance. If we try many different specifications and pick the one that fits the expected or desired outcome most closely, we will overstate the significance of those results. This is specification search bias (or data-mining bias or data-snooping bias). To correct this bias, when we try different regression specifications and pick the "best" one, we should, in principle, adjust significance levels and possibly coefficient estimates. For this particular example an appropriate correction, the Bonferroni correction, is known.[38]

However, for most types of specification searching the appropriate correction to significance levels is not known. For example, different explanatory variables could be tried. Maybe interest rates could be included in the regression, maybe we could use provincial GDP as an explanatory variable, maybe we could use land rent (a cost for dealers), etc. Another type of specification searching involves trying different time periods in the analysis, using or not using observations from after the class period, or possibly using lagged variables. Or we might try different functional forms. These are often important steps in finding the best specification but, unfortunately, for these and most other types of specification searching, the appropriate correction to significance levels is either very difficult or impossible to determine in precise form.

A standard recommendation is to divide the data into two parts – one part used to select the best specification and the other used for actual estimation. Then the results from the second estimation are sometimes thought to be free of specification search (data-mining) bias. However, dividing the sample into two parts may not be feasible in price-overcharge situations. Even if it is feasible, because only part of the data is used for estimation, the estimated parameters are estimated less precisely than if the full sample is used (i.e. the significance level may be correctly identified but the coefficient is inaccurately estimated). This process is also subject to other problems.[39]

The other standard recommendation, which is highly relevant for price-overcharge estimation, is to do sensitivity analysis, which means trying and reporting different specifications to see how much the results are changed by various changes in specification or procedure. If we use many reasonable but different specifications and consistently find similar results, that greatly strengthens our confidence in those results. Sensitivity analysis has the opposite effect of specification search bias.

In addition, if the apparent p-values and standard errors are very low and the confidence intervals are very tight then, even if we did some specification searching, it is still likely that the coefficient estimates are meaningful. In such a case, even if we knew the proper adjustments to account for specification searching and made them, we would likely still find significant and meaningful results, although not quite as significant as they might seem at face value.

We believe that there are two lessons to be learned from consideration of specification search bias. The first is that the stated confidence intervals and significance levels produced by statistical software should not be taken literally. A 95% confidence interval does not really mean that there is a 95% chance that a coefficient lies in that interval. If the stated result is the outcome of even modest amounts of specification searching, then the stated confidence level is too high.[40] The second lesson is that statistical analysis is very valuable, but assessing its value is an art as well as a science and requires judgment.

In assessing statistical work, one important characteristic to look for is consistency. For example, in pass-through estimation, if the results of statistical analysis from many different vendors are consistent with each other and are consistent with market structure evidence and with statements of industry participants and with prior analysis of pass-through in similar situations, then the result should be taken very seriously. However, an isolated and surprising result should be regarded with much more caution.

VII. Concluding Remarks

This article identifies what we view as important principles in estimating damages arising from price-fixing and related anti-competitive actions. We focus particularly on damage estimation for indirect purchasers, although much of the material we cover is relevant to direct purchasers as well.

With respect to indirect purchasers we recognize that the damages they face depend on three components: the at-issue revenues of cartel members, the proportional overcharge imposed by the cartel and the degree of pass-through from direct to indirect purchasers.

In the paper we review estimation of all three of these components, although our discussion of revenue is brief. We have an extensive discussion of estimating overcharges. Sometimes we estimate the nominal overcharge, which is the actual price minus the but-for price, and sometimes we estimate the proportional overcharge – the fraction of the actual price that is an overcharge. We present five general methods, although only two are now widely used in our experience – reduced form estimation of prices and use of comparator benchmarks. It is also possible to use these two methods in combination.

We also provide a detailed discussion of pass-through, which is very important for indirect purchaser cases. We point out a pass-through rate of 100% is of particular interest as it is what we expect if the downstream sector is highly competitive and has constant average cost at the industry level.

The other major part of the paper tries to open up the black box of econometrics, at least to some extent. Regression analysis is a particularly important econometric tool. We provide a brief overview of how regression results are normally reported, emphasizing the role of statistical significance and related concepts. We also discuss the interpretation of regression results, placing particular emphasis on the role of specification searching (sometimes called *data-mining*).

Overall, we believe that a great deal of progress in damage estimation and related topics has been made in the past two decades. In addition, data availability has significantly improved and computing power has increased greatly. Therefore, good estimates of damages from price-fixing and related anticompetitive practices can often be obtained.

ENDNOTES

* Brander: Asia Pacific Professor in International Business and Public Policy, Sauder School of Business, University of British Columbia. Ross: UPS Foundation Professor of Regulation and Competition Policy, Sauder School of Business, University of British Columbia. We thank Jennifer Ng for very valuable comments and other help on this paper. In addition, we would like to acknowledge the important contributions made by several of our colleagues to our understanding of the issues discussed in this paper, including J. J. Camp, Tom Davidoff, David Jones, Reidar Mogerman, Janet Netz, Dov Rothman and Kevin Wright, though the views expressed here are our own and may not reflect their views. Financial support was provided by the Phelps Centre for the Study of Government and Business at the Sauder School of Business.

[1] A list of private actions in price-fixing cases is provided in John M Connor, "Private Recoveries in International Cartel Cases Worldwide: What Do the Data Show?" (2012) American Antitrust Institute Working Paper No 12-03, online: <http://ssrn.com/abstract=2165431>. The Canadian Bar Association maintains a database on class actions (including those related to price-fixing) ["Class Action Database", online: The Canadian Bar Association <http://www.cba.org/Publications-Resources/Class-Action-Database>].

[2] *Sun-Rype Products Ltd v Archer Daniels Midland Company*, 2013 SCC 58, [2013] 3 SCR 545; *Pro-Sys Consultants Ltd v Microsoft Corporation*, 2013 SCC 57, [2013] 3 SCR 477; *Infineon Technologies AG v Option Consommateurs*, 2013 SCC 59, [2013] 3 SCR 600.

[3] The landmark U.S. Supreme Court case establishing the direct purchaser rule for U.S. federal cases is *Illinois Brick Co v Illinois*, 431 US 720 (1977).

[4] James A Brander & Thomas W Ross, "Estimating Damages from Price-Fixing" (2006) 3:1 *Can Class Action Rev* 335.

[5] Sales taxes paid by purchasers are not normally included in damages. So the relevant price paid by the purchasers is the pre-tax price, which is the same price received by the seller in direct purchaser actions.

[6] See, for example, Brander & Ross, *supra* note 4 and Leonardo J Basso & Thomas W Ross, "Measuring the True Harm from Price-Fixing to Both Direct and Indirect Purchasers" (2010) 58:4 *J Ind Econ* 895.

[7] We are not taking the position that we should never try to measure harm not captured by the overcharges on units still purchased – rather here we are choosing to focus our discussion on the measure of damages most commonly applied in practice. It is important to acknowledge, however, that under some circumstances this measure can significantly understate the real harm caused to direct and indirect purchasers by price-fixing. On this see, for example, Basso & Ross, *supra* note 6.

[8] Sometimes the proportional overcharge is defined relative to the but-for price so it is $(p^A - p^B)/p^B$. This has the disadvantage that the proportional overcharge can be a very large number, approaching infinity as the actual price becomes large relative to the but-for price and it is often less convenient for damage calculations.

[9] The algebra with two product categories is as follows. We denote product categories using the subscripts 1 and 2. $D=(p_1^A - p_1^B)\, q_1^A + (p_2^A - p_2^B)\, q_2^A = v_1R_1 + v_2R_2 = (v_1R_1 + v_2R_2)[(R_1 + R_2)/(R_1 + R_2)] = (v_1w_1 + v_2 w_2)R = vR$, where $v = v_1w_1 + v_2w_2$ and $w_i = R_i/(R_1 + R_2)$ for $i = 1, 2$. Also, $R = R_1 + R_2 = p_1^A q_1^A + p_2^A q_2^A$. Adding more terms (more product categories) works in exactly the same way.

[10] This analysis focuses only on damages arising from sales made by cartel members. It is possible that other sellers (non-members of the cartel) may exist and may sell at higher prices than they otherwise would due to the actions of the cartel (the "umbrella" theory). Buyers of those products may also suffer damage. We do not comment on the status of such buyers or such damages.

[11] In its 2013 decisions on class actions, the Supreme Court of Canada

recognized the possible conflict between different categories of class members over the distribution of aggregate damages. However, the Court ruled that this conflict is not an impediment to certification and is properly dealt with as distribution issue later in the process.

[12] We note in passing that plaintiffs in indirect purchaser actions may choose to make a claim based on "unjust enrichment" in addition to (or instead of) a claim based on economic harm to plaintiffs. Unjust enrichment can be estimated as $D = vR$, where R is the at-issue revenue received by the cartel. This formula would apply regardless of the extent of pass-through as it shows the extra revenue obtained by defendants arising from overcharges on the quantity sold. In its 2013 decisions on class actions, the Supreme Court did not rule out unjust enrichment claims.

[13] See, in particular, John M Connor, "Cartel Overcharges" (2014) 26, ed by James Langenfeld, *L Econ Class Actions* (*Research in Law and Economics*) 249. This paper surveys over 700 studies, providing over 2000 overcharge estimates.

[14] In addition, post-conspiracy prices might be higher than in the but-for world if the cartel firms are able to achieve some form of tacit collusion based on their cartel experience. And pre-conspiracy "price-war" periods might have lower prices than would be observed in a normal but-for situation.

[15] From the Hynix plea agreement in the U.S. DRAM action: *United States of America v Hynix Semiconductor Inc*, (ND Cal 2005), online: <http://www.usdoj.gov/atr/cases/f209200/209231.pdf> "At certain times during the relevant period, DRAM prices decreased significantly. Nevertheless, the Defendant and its coconspirators reached agreements to limit the rate of price declines, which were achieved with varying levels of effectiveness."

[16] White (2001) suggests that this was John Connor's approach in his work for plaintiffs in the U.S. lysine damages case, but Connor might not see it that way. Connor (2001) characterizes his approach as employing the "before and after" method. In fact, in this case is hard to distinguish the approaches here because the price before approached levels consistent with reasonable estimates of average cost. See Lawrence J White, "Lysine and Price Fixing: How Long? How Severe?" (2001) 18:1 *Rev Ind Org*, 23 and John M Connor, " 'Our Customers are Our Enemies': The Lysine Cartel of 1992-1995" (2001) 18:1 *Rev Ind Org* 5. Hastings and Williams (2016) describe a similar approach, from a tied-selling case in which a competitive but-for price for cable television set-top boxes (that the defendant had tied to its sale of premium television services) was constructed from data on the defendant's direct costs and an assessment of its opportunity cost of capital. Adding in the opportunity cost of capital generates a but-for price that allows for a competitive rate of return. Justine S Hastings & Michael W Williams, "What is a 'But-For' World" (2016) 31:1 *Antitrust* 102.

[17] These economic models are described in standard textbooks on industrial organization such as Dennis W Carlton & Jeffrey M Perloff, *Modern Industrial Organization*, 4th ed (Boston: Pearson, 2005).

[18] For example, since prices will depend on the nature of competition between

firms (as well as cost and demand factors) a structural approach will require the analyst to adopt a specific model of competition.

[19] For a good discussion of the issues see the following articles: Joshua D Angrist & Jörn-Steffen Pischke, "The Credibility Revolution in Empirical Economics: How Better Research Design is Taking the Con out of Econometrics" (2010) 24:2 *J Econ Perspect* 3; James F Nieberding, "Estimating Overcharges in Antitrust Cases Using a Reduced-Form Approach: Methods and Issues" (2006) 9:2 *J Appl Econ* 361, and Douglas J Zona, "Structural Approaches to Estimating Overcharges in Price-Fixing Cases" (2011) 77:2 *Antitrust LJ* 473.

[20] To be precise, equation (4) is meant to hold at any given point in time – so that, for example, the price at some particular point in time will be explained by the function with the right-hand side variables' values also taken at that same point in time. It is common for economists to put subscripts on the variables indicating which period the data is drawn from, but we omit them here to reduce clutter.

[21] The term we prefer here, comparator benchmarks, was used in Oxera Consulting Ltd et al, *Quantifying antitrust damages: Towards non-binding guidance for courts* (Luxembourg: Publications Office of the European Union, 2009), online: <http://www.oxera.com/Oxera/media/Oxera/Quantifying-antitrust-damages.pdf?ext=.pdf>.

[22] Justin McCrary & Daniel L Rubinfeld, "Measuring benchmark damages in antitrust litigation" (2014) 3:1 *J Econ Methods* 63.

[23] See 2014 BCSC 532, varied 2015 BCCA 363.

[24] *Pro-Sys Consultants v Microsoft Corporation*, 2010 BCSC 285, [2010] BCJ No 380 (QL) at 25.

[25] A basic description of this case is set out in *Steele v Toyota Canada Inc*, 2008 BCSC 1063, 295 DLR (4th) 653 . Approval of the Settlement Agreement can be found at Steele v Toyota Canada Inc, 2015 BCSC 1014.

[26] Under the Access program, Toyota organized a system under which individual dealers would suggest a reasonable "firm price", normally below the list price and, after collating the suggestions, Toyota would inform the dealers of the "Access price". Dealers were expected (but not strictly required) to charge this price as a firm price. The system was advertised to consumers as providing a "no hassle" price under which everyone would get a fair price instead of the final price being dependent on bargaining effort and skills.

[27] If we use as a benchmark prices from other producers who are not part of the cartel but are nevertheless in the market, we should consider the possibility that the prices of non-cartel firms might be influenced upward by the higher cartel prices. This is the so-called umbrella effect. In such a case, these benchmark firms' prices would be expected to be above true but-for prices and using them could lead to an underestimate of the real overcharge.

[28] See, for example, Basso & Ross, *supra* note 6.

[29] The elasticity of demand measures the sensitivity of quantity demanded to price changes. Specifically, it is the percentage change in quantity demanded

caused by a 1% increase in price. When the demand elasticity itself does not change as price changes, firms will set constant margins. For example, if the monopolist direct purchaser has a constant demand elasticity equal to 5, it will set a profit margin of 20% to maximize its profits. To maintain that 20% margin, any cost increase of, say, $1 must lead to a downstream price increase of $1.20

[30] For proofs of these statements and a more complete treatment of pass-through rates in the context of oligopolies, see Basso & Ross, *supra* note 6.

[31] This approach can still require adjustments necessitated by delays in price adjustments downstream.

[32] To be clear, this is not to say that an effect is absent, it can simply be hard to measure with precision. Theoretically this can be overcome if we have enough high quality data.

[33] Changing prices carries a cost, for example, in terms of adjusting price lists, labels and advertising and the risk of customer annoyance and resistance. When the price-fixed component is a relatively small part of the total costs of the downstream firm, it will make sense to delay adjusting prices until other cost or demand changes have made a non-trivial adjustment profitable. When the price adjustment is made, however, it should be expected to take into consideration all the firm's costs including those subject to price-fixing.

[34] One very good general introductory textbook on econometrics containing many applied examples (although no damage estimation examples) is Jeffrey M Wooldridge, *Introductory Econometrics: A Modern Approach*, 6th ed (Boston: Cengage Learning, 2015).

[35] See, e.g. Daniel L Rubinfeld, "Econometric Issues in Antitrust Analysis" (2010) 166:1 *J Inst Theor Econ* 62; Jonathan B Baker & Daniel L Rubinfeld, "Empirical Methods in Antitrust: Review and Evidence" (1999) 1:1 *Am L Econ Rev*, 386; Rober E Hall & Victoria A Lazear, "Reference Guide on Estimation of Economic Losses in Damages Awards" (2000) *Reference Manual on Scientific Evidence*, online: <http://www.fjc.gov/public/pdf.nsf/lookup/12.econ_loss.pdf/$File/12.econ_loss.pdf>; and Oxera Consulting Ltd et al, *supra* note 21.

[36] These measures relate to classical or frequentist methods. It is also possible to report related measures based on Bayesian methods, but that is rare in our experience of damage estimation and related areas and we will not discuss Bayesian methods here.

[37] It is rare that the law of one price holds exactly, More generally, significant positive estimates for a_1 and a_2 imply some level of integration between the markets.

[38] See, for example, Wolfram MathWorld, "Bonferroni Correction", online: <http://mathworld.wolfram.com/BonferroniCorrection.html>.

[39] Atsushi Inoue & Lutz Kilian, "In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?" (2005) 23:4 Economet Rev 371.

[40] Even if we are confident that we have the correct specification without doing any specification searching, the statement just made is not quite right. Strictly

speaking, we should say that if we did repeated sampling, the 95% confidence interval would include the true value 95% of the time.