# Testing Competing Nash-in-Nash Bargaining Models

VERY PRELIMINARY AND INCOMPLETE, PLEASE DO NOT CITE OR CIRCULATE

João Granja
Vladimir Pavlov

March 18, 2024

**Abstract**

Specifying empirical Nash-in-Nash bargaining models requires the analyst to make a number of modeling choices, such as the level at which parties negotiate and how disagreement payoffs are determined. The empirical bargaining literature features different combinations of these choices. We ask whether these auxiliary assumptions are consequential for the predictions of these models, and whether data can distinguish between them. We extend recently proposed tests of firm pricing behavior to test the competing Nash-in-Nash models defined by different combinations of these auxiliary assumptions. We apply this test to data on wholesale prices in the Washington state legal cannabis industry. A main motivation for the legalization of cannabis in Washington was the revenue obtained from excise taxes. To illustrate the economic significance of employing the most appropriate bargaining model, we compute the Laffer curves implied by the different models.

Granja: Department of Economics, University College London, joao.granja@ucl.ac.uk.
Pavlov: School of Management, University College London, v.pavlov@ucl.ac.uk.
We thank Kirill Borusyak, Lorenzo Magnolfi, Lars Nesheim, and Gabriel Ulyssea for helpful conversations.

# 1 Introduction

Many markets feature negotiated prices. Empirical IO and other economists have widely adopted the Nash-in-Nash framework of Horn and Wolinsky (1988) as the basis for empirical models of negotiated contract terms. Since the first applications by Crawford and Yurukoglu (2012) and Grennan (2013), this framework has become increasingly popular. Applications in IO abound. To cite just a few: Gowrisankaran, Nevo, and Town (2015), Ho and Lee (2017), Crawford, Lee, Whinston, and Yurukoglu (2018), De los Santos, O'Brien, and Wildenbeest (2021), Ellickson, Lovett, Sunada, and Kong (2022). Applications of the Nash-in-Nash framework are not limited to IO. Applications in International Trade include Bagwell, Staiger, and Yurukoglu (2021) and Alviarez, Fioretti, Kikkawa, and Morlacco (2023) and there are many applications to Labor Economics.

A set of contracts is a Nash-in-Nash equilibrium if each contract maximizes a Nash product given the other contracts. To operationalize this concept, the analyst has to make many modeling choices. First, at what level do negotiations occur? At the firm level? At the product level? Different papers make different assumptions. Gowrisankaran et al. (2015) and Crawford and Yurukoglu (2012) make the former assumption, whereas Grennan (2013), Crawford et al. (2018), Ellickson et al. (2022) make the latter.[1] Second, how should disagreement payoffs be specified? Do firms adjust other actions following disagreement? If so, which firms? Again, different papers make different assumptions. Draganska, Klapper, and Villas-Boas (2010), Ho and Lee (2017), Crawford et al. (2018), and De los Santos et al. (2021) assume that no actions are changed following a disagreement. In contrast, in the context of negotiations between manufacturers and retailers, Crawford and Yurukoglu (2012) and Ellickson et al. (2022) assume that downstream firms adjust their remaining prices when a negotiation breaks down.

This diversity of approaches raises two questions. First, do these approaches lead to different predictions regarding negotiated contract terms? Second, can data distinguish between the different models implied by these different modeling choices?

These are the questions we tackle in this paper. We start by setting a general Nash-in-Nash framework that accommodates all of the choices discussed above. We then show how to extend the testing framework of Backus, Conlon, and Sinkinson (2021) and Duarte, Magnolfi, Sølvsten, and Sullivan (2023) to

1

test the competing bargaining models implied by different choices of the level of negotiations and firms' behavior following disagreement.

We apply these methods to data from the legal cannabis industry in the state of Washington. That industry contains three vertically related types of firms: producers, processors, and retailers. The latter cannot vertically integrate, and therefore must acquire the goods that they sell from processors. We test which of the multiple Nash-in-Nash models provides the best description of the prices at which retailers and processors trade. To show that distinguishing between these different models is economically important, we compute the Laffer curves implied by the different models of bargaining, and show that different models have different implications for the revenue-maximizing level of excise taxes. This is a contribution in its own right, as it extends the recent literature on taxation with market power, such as Miravete, Seim, and Thurk (2018) and Hollenbeck and Uetake (2021) to bargaining models.[2]

The rest of the paper is organized as follows. Section 2 introduces a Nash-in-Nash framework that allows for different levels of negotiation and disagreement payoffs. Section 3 discusses the Washington state legal cannabis industry and the data. Section 4 leverages existing theoretical results on how the predictions of bargaining models depend on auxiliary assumptions regarding disagreement payoffs and exogenous variation in the number of retailers due to a license lottery to test those theoretical predictions and thus shed some light on which specification of disagreement payoffs may be best fit our data. Section 5 discusses the extension of the Backus et al. (2021) and Duarte et al. (2023) tests to Nash-in-Nash models. Section 6 concludes.

# 2 Model

We consider a market with $R$ retailers indexed by $r = 1, \dots, R$. A *product* is a combination of observed physical attributes, a manufacturer, and a retailer. There are $J$ products indexed by $j = 1, \dots, J$ and we let $\mathcal{J} := \{1, \dots, J\}$ denote the set of all products in the market. We denote by $\mathcal{J}_r \subseteq \mathcal{J}$ the set of products sold by retailer $r$. There are also $M$ manufacturers indexed by $m = 1, \dots, M$. We let $\mathcal{J}_m \subseteq \mathcal{J}$ be the set of products manufactured by firm $m$. Note that given our definition of a product, the set $\mathcal{J}_m$ can contain the same physical product sold to multiple retailers. Finally, we let $\mathcal{J}_{rm} := \mathcal{J}_r \cap \mathcal{J}_m$ be the set of products

[2]This exercise is work in progress.

sold by retailer $r$ and produced by manufacturer $m$.

Here we take as given market demand functions that depend on prices and the available products: $D_j(\boldsymbol{p}; \mathcal{J})$, $j = 1, \ldots, J$. Elsewhere we derive these demand functions from a model of consumer choices. The profits of a retailer given a retail price vector $\boldsymbol{p}$, its vector of wholesale prices $\boldsymbol{w}_r$ and the set of products in the market $\mathcal{J}$ are given by

$$\pi^r(\boldsymbol{p}, \boldsymbol{w}_r; \mathcal{J}) = \sum_{j \in \mathcal{J}_r} (p_j - w_j) D_j(\boldsymbol{p}; \mathcal{J}) \tag{1}$$

Retailers compete à la Bertrand. Denote the equilibrium prices by $\boldsymbol{p}^*(\boldsymbol{w}; \mathcal{J})$. We then have quantities and retailers' profits as a function of wholesale prices: $D_j^*(\boldsymbol{w}; \mathcal{J}) := D_j(\boldsymbol{p}^*(\boldsymbol{w}; \mathcal{J}); \mathcal{J})$ and $\pi^r(\boldsymbol{w}; \mathcal{J}) := \pi^r(\boldsymbol{p}^*(\boldsymbol{w}; \mathcal{J}), \boldsymbol{w}_r; \mathcal{J})$, respectively.

Manufacturers' profits are given by

$$\pi^m(\boldsymbol{w}, \boldsymbol{c}_m; \mathcal{J}) = \sum_{r=1}^{R} \sum_{j \in \mathcal{J}_{rm}} (w_j - c_j) D_j^*(\boldsymbol{w}; \mathcal{J}) \ , \tag{2}$$

where $c_j$ is the marginal cost of producing good $j$ and $\boldsymbol{c}_m$ is the vector of manufacturer $m$'s marginal costs of producing the goods in $\mathcal{J}_m$.

## 2.1 Bargaining

Firms bargain over the linear wholesale prices $w_j$. We assume that the wholesale price vector is determined by Nash-in-Nash bargaining. Specifically, suppose that for each retailer-manufacturer pair $(r, m)$ we have a partition $\{\mathcal{J}_{rm}^n\}_{n=1}^{N_{rm}}$ of $\mathcal{J}_{rm}$. Here, $n$ indexes different negotiations between manufacturer $m$ and retailer $r$.[3] Given prices for all other negotiations, $\boldsymbol{w}_{-\mathcal{J}_{rm}^n}$, define the manufacturer and retailer surplus from their $n$-th negotiation respectively by

$$\Delta\pi^m(\hat{\boldsymbol{w}}, \boldsymbol{w}_{-\mathcal{J}_{rm}^n}, \boldsymbol{c}_m, \mathcal{J}, \mathcal{J}_{rm}^n) := \pi^m(\hat{\boldsymbol{w}}, \boldsymbol{w}_{-\mathcal{J}_{rm}^n}; \boldsymbol{c}_m, \mathcal{J}) - \pi_D^m \tag{3}$$

and

$$\Delta\pi^r(\hat{\boldsymbol{w}}, \boldsymbol{w}_{-\mathcal{J}_{rm}^n}, \mathcal{J}, \mathcal{J}_{rm}^n) := \pi^r(\hat{\boldsymbol{w}}, \boldsymbol{w}_{-\mathcal{J}_{rm}^n}; \mathcal{J}) - \pi_D^r \ , \tag{4}$$

[3] We introduce this notion of a partition because it allows us to consider the cases where $S = 1$ (firms bargain over all products at once), the case $S = |\mathcal{J}_{rm}|$ (firms bargain separately product by product, as in Grennan (2013)), and all intermediate cases (e.g., separate bargaining over the groups of goods associated with different brands of a manufacturer).

where $\pi_D^m$ and $\pi_D^r$ are the disagreement payoffs of the manufacturer and the retailer, respectively. Then, the vector of equilibrium wholesale prices $\boldsymbol{w}^*$ is such that

$$\boldsymbol{w}^*_{\mathcal{J}^n_{rm}} \in \underset{\hat{\boldsymbol{w}} \in \mathbb{R}^{|\mathcal{J}^n_{rm}|}}{\arg\max} \, [\Delta\pi^m(\hat{\boldsymbol{w}}, \boldsymbol{w}^*, \boldsymbol{c}_m, \mathcal{J}, \mathcal{J}^n_{rm})]^\beta \times [\Delta\pi^r(\hat{\boldsymbol{w}}, \boldsymbol{w}^*, \mathcal{J}, \mathcal{J}^n_{rm})]^{1-\beta} \tag{5}$$

In equations (3) and (4) we have been purposely vague about the disagreement points $\pi_D^m$ and $\pi_D^r$. There are multiple, equally reasonable, ways to model disagreement payoffs. Quoting from Horn and Wolinsky (1988):

> A proper specification of the [disagreement point] is not an obvious matter. (...) this choice [the model they adopt] would correspond to a situation in which, when firm $i$ and the supplier cannot agree, the firm earns zero profit and the other firm, $j$, operates at the anticipated equilibrium. It should be noted, however, that there are other plausible specifications. For example, if as long as firm $i$ does not reach agreement with the [single] supplier firm $j$ can act as a monopoly, then the disagreement position of the supplier should be modified...

Here we consider the two alternatives outlined by Horn and Wolinsky (1988) and others. The two models outlined in the quote above correspond to the two cases in Iozzi and Valletti (2014): unobservable breakdowns and observable breakdowns, respectively. These cases suffice for the aforementioned papers because they deal with single-product retailers. We consider multi-product retailers. In this case, another reasonable alternative arises: the retailer whose negotiation fails observes that fact and adjusts its price; the other retailers do not. We will refer to these three cases as *No Repricing*, *Full Repricing*, and *Partial Repricing*. We denote the corresponding disagreement profits by $\pi^\ell_{NR}, \pi^\ell_{FR}$ and $\pi^\ell_{PR}$, where $\ell$ indexes either a manufacturer or a retailer.

The disagreement payoffs in these cases are as follows:

- No Repricing

$$\pi^r_{NR}(\boldsymbol{w}^*, \mathcal{J}, \mathcal{J}^n_{rm}) = \sum_{j \in \mathcal{J}_r \setminus \mathcal{J}^n_{rm}} \left[p^*_j(\boldsymbol{w}^*; \mathcal{J}) - w^*_j\right] D_j(\boldsymbol{p}^*(\boldsymbol{w}^*; \mathcal{J}); \mathcal{J} \setminus \mathcal{J}^n_{rm})$$

$$\pi^m_{NR}(\boldsymbol{w}^*, \boldsymbol{c}_m, \mathcal{J}, \mathcal{J}^n_{rm}) = \sum_{r'=1}^{R} \sum_{j \in \mathcal{J}_{r'm} \setminus \mathcal{J}^n_{rm}} (w^*_j - c_j) D_j(\boldsymbol{p}^*(\boldsymbol{w}^*; \mathcal{J}); \mathcal{J} \setminus \mathcal{J}^n_{rm})$$

The retailer's payoff is equal to the sum of its variable profits arising from each of the products it sells except for those that are part of its $n$-th negotiation with manufacturer $m$. The variable profits are computed using the retail prices that would prevail in the downstream Nash-Bertrand equilibrium given the equilibrium wholesale prices $w^*$. Note that consumer choices do adjust to the absence of the products in $\mathcal{J}^n_{rm}$.

- Full repricing

  As above, but the prices are the Nash-Bertrand equilibrium prices given the reduced set of products $\mathcal{J} \setminus \mathcal{J}^n_{rm}$ and the equilibrium wholesale prices of these products, $\boldsymbol{w}^*_{-\mathcal{J}^n_{rm}}$, i.e., $\boldsymbol{p}^*(\boldsymbol{w}^*_{-\mathcal{J}^n_{rm}}; \mathcal{J} \setminus \mathcal{J}^n_{rm})$. As before, demand depends on the reduced set of products $\mathcal{J} \setminus \mathcal{J}^n_{rm}$.

- Partial repricing

  As above, but only the focal retailer changes its prices. Retailers other than $r$ continue to set $\boldsymbol{p}^*_{-r}(\boldsymbol{w}^*; \mathcal{J})$. Retailer $r$ instead sets its best response to these prices given its reduced product portfolio, i.e., $\boldsymbol{p}^{BR}(\boldsymbol{p}^*_{-r}(\boldsymbol{w}^*; \mathcal{J}); \mathcal{J} \setminus \mathcal{J}^n_{rm})$.

### 2.1.1 The Economics Behind the Alternative Models

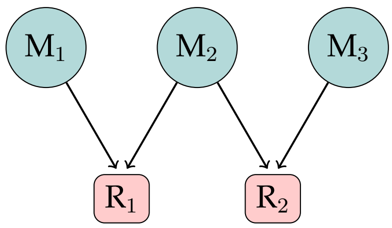


Figure 1: A supply-chain network

Consider the supply chain network in figure 1. We denote a product (i.e., a retailer-manufacturer pair) by the ordered pair $(r, m)$; we denote its price by $p^r_m$. Suppose that the relationship between manufacturer 2 and retailer 1 breaks down. What adjustments ensue? First, there is a direct effect: $R_1$ offers less variety, and some consumers will switch to $R_2$. Some of these consumers will switch to good $(2, 2)$, but in general $M_2$ will experience a reduction in the quantity it sells. Due to its reduced product portfolio, retailer 1 will reduce the price of $M_1$'s product, $p^1_1$. This will make some consumers switch from $R_2$ to $R_1$, further reducing the quantity sold by $M_2$. We conclude that $R_1$'s ability to adjust its prices increases $M_2$'s relationship surplus and thus $R_1$'s bargaining leverage. We thus expect that the ability to reprice will lead to lower wholesale mark-ups.[4]

Next, what happens if $R_2$ can also adjust its prices? Here, note that there are two counteracting effects. On the one hand, $R_1$ is a weaker competitor, because it only carries one product. On the other hand, $R_1$ has reduced the price of good $(1, 1)$, making it a stronger competitor. It is a priori unclear which effect dominates. If the portfolio effect dominates, $R_2$ will increase its prices. This will further reduce the quantity sold by $M_2$. Therefore, when all firms can adjust their prices following a breakdown and the portfolio effect dominates, $M_2$'s relationship surplus is even larger than when only the focal retailer reprices, thus leading to even lower wholesale markups. If the pricing effect dominates, $R_2$ will reduce its prices, thus increasing the quantity sold by $M_2$. This makes $M_2$'s relationship surplus *smaller* than when only the focal retailer can reprice, thus leading to larger wholesale mark-ups.[5]

It is clear (but should be spelled out in greater detail) that the effects above will depend on how substitutable manufacturers' products and retailers are. This thus suggests that as we vary market structure (the number of retailers, their ownership, and their product portfolios) and the nature of demand (consumers' demographic characteristics), the different bargaining models above will generate different comparative statics predictions. These are comparative statics whose counterparts are present in our data. That is the variation in the data that will allow us to tell the competing models apart.

### 2.1.2 The System of NiN FOCs

Take the logarithm of (5) and differentiate to obtain

$$\beta \frac{\partial_{w_j} \Delta \pi^m}{\Delta \pi^m} + (1 - \beta) \frac{\partial_{w_j} \Delta \pi^r}{\Delta \pi^r} = 0 \tag{6}$$

Multiply both sides by $\Delta \pi^m \Delta \pi^r$ to obtain

$$\beta \Delta \pi^r \partial_{w_j} \Delta \pi^m + (1 - \beta) \Delta \pi^m \partial_{w_j} \Delta \pi^r = 0 \tag{7}$$

[4]This effect is strongest when the residual demand faced by the retailer is more elastic, i.e., when it faces tougher competition. We thus conclude that the ability to reprice should lead to lower predicted upstream mark-ups, and more so the tougher the competition in the downstream market.

[5]This effect will again be strongest the more elastic retailers' residual demands are, i.e., the more competition they face.

Next note that

$$\partial_{w_j}\Delta\pi^m = \sum_{r=1}^{R}\sum_{k\in\mathcal{J}_{rm}}(w_k - c_k)\partial_{w_j}D_k^*(\boldsymbol{w}^*;\mathcal{J}) + D_j^*(\boldsymbol{w}^*;\mathcal{J}) \tag{8}$$

and

$$\partial_{w_j}\Delta\pi^r = -D^*(\boldsymbol{w}^*;\mathcal{J}) \;, \tag{9}$$

where the last equation is due to the Envelope Theorem. Plugging these results into equation (7) gives

$$\begin{aligned}
0 = {} & \beta\sum_{r=1}^{R}\sum_{k\in\mathcal{J}_{rm}}(w_k - c_k)\partial_{w_j}D_k^*(\boldsymbol{w}^*;\mathcal{J})\Delta\pi^r \\
& + \beta D_j^*(\boldsymbol{w}^*;\mathcal{J})\Delta\pi^r \\
& - (1-\beta)D_j^*(\boldsymbol{w}^*;\mathcal{J})\left[\sum_{r=1}^{R}\sum_{j\in\mathcal{J}_{rm}}(w_j - c_j)D_j^*(\boldsymbol{w}^*;\mathcal{J}) - \sum_{r=1}^{R}\sum_{j\in\tilde{\mathcal{J}}_{rm}}(w_j - c_j)\tilde{D}_j\right]
\end{aligned} \tag{10}$$

where we have used that

$$\Delta\pi^m = \sum_{r=1}^{R}\sum_{j\in\mathcal{J}_{rm}}(w_j - c_j)D_j^*(\boldsymbol{w}^*;\mathcal{J}) - \sum_{r=1}^{R}\sum_{j\in\tilde{\mathcal{J}}_{rm}}(w_j - c_j)\tilde{D}_j \;, \tag{11}$$

where $\tilde{\mathcal{J}}_{rm}$ denotes the set of products that $r$ and $m$ negotiate over after disagreement in the specific negotiation under consideration and $\tilde{D}_j$ is the demand for product $j$ after disagreement in the negotiation under consideration.

Now divide equation (10) by $\beta\Delta\pi^r$, subtract $D_j^*(\boldsymbol{w}^*;\mathcal{J})$ from both sides, and collect summation terms to obtain

$$-D_j^*(\boldsymbol{w}^*;\mathcal{J}) = \sum_{r=1}^{R}\sum_{k\in\mathcal{J}_{rm}}\left\{\partial_{w_j}D_k^*(\boldsymbol{w}^*;\mathcal{J}) - B_{jk}(\beta,\boldsymbol{w}^*,\mathcal{J},\Delta\pi^r)\right\}(w_k - c_k) \tag{12}$$

where

$$B_{jk}(\beta,\boldsymbol{w}^*,\mathcal{J},\Delta\pi^r) = \frac{1-\beta}{\beta}\frac{D_j^*(\boldsymbol{w}^*;\mathcal{J})}{\Delta\pi^r}\left[D_k^*(\boldsymbol{w}^*;\mathcal{J}) - \mathbf{1}\{k\notin\mathcal{J}_{rm}^n\}\tilde{D}_k\right] \;. \tag{13}$$

Stacking these equations across products $j\in\mathcal{J}$ yields

**Proposition 1.** *If $\boldsymbol{w}^*$ is a Nash-in-Nash equilibrium, then*

$$\left\{[D_w\sigma^*(\boldsymbol{w}^*;\mathcal{J}) - B(\boldsymbol{\beta},\boldsymbol{w}^*;\mathcal{J})]\odot\Omega\right\}(\boldsymbol{w}^* - \boldsymbol{c}) = -D^*(\boldsymbol{w}^*;\mathcal{J}) \tag{14}$$

*where* $D^*(\boldsymbol{w}^*; \mathcal{J}) = D(\boldsymbol{p}^*(\boldsymbol{w}^*; \mathcal{J}); \mathcal{J})$ *and*

$$B(\boldsymbol{\beta}, \boldsymbol{w}^*; \mathcal{J}) = \left[\frac{(1-\beta)}{\beta}\frac{D_j^*(\boldsymbol{w}^*; \mathcal{J})}{\Delta\pi^r}\left[D_k^*(\boldsymbol{w}^*; \mathcal{J}) - \mathbf{1}\{k \notin J_{rm}^n\}\tilde{D}_k\right]\right]_{jk}$$

# 3 Institutional Setting and Data

In November 2012, the state of Washington legalized recreational cannabis. The first licenses to produce and process cannabis were issued in March 2014. Retail licenses wewre first issued in July 2014. The Washington State Liquor and Cannabis Board (the sector regulator, henceforth WLCB) imposed a state-wide cap on the number of retail licenses. The number of retail licenses was first limited to 334, and that number was later raised to 556. Retail licenses were allocated to counties and then to jurisdictions within counties. There were two types of retail licenses, called *designated* and *at-large*. Designated licenses apply to a specific location, typically a large city such as Seattle or Tacoma. At-large licenses can be used to set up a dispensary anywhere in the county other than designated locations.

Potential retail licensees had an opportunity to apply for licenses specifying their jurisdiction of interest (either a county or one of those covered by the designated licenses). This process led to oversubscribed jurisdictions. To allocate licenses to retailers in oversubscribed jurisdictions, the WLCB conducted lotteries. The lotteries generate exogenous variation in the number of retailers in certain jurisdictions. This will be useful for the econometric analysis that follows. The state of Washington is a closed market: importing and exporting cannabis products is illegal. Finally, we note that retailers cannot vertically integrate with processors nor producers. Processors and producers can, and do, vertically integrate.

The WLCB conducts very detailed data on the market. From the WLCB we obtained data on wholesale and retail transactions, as well as the lottery. The transaction data records the type of the product (e.g., usable cannabis or edibles), the strain of cannabis, the quantity transacted, and the price paid. In the case of wholesale transactions, we know the identity of the parties involved. For retail transactions, we observe the identity of the retailer, but cannot track consumers. The retail lottery data includes all applicants in a jurisdiction, the location linked to the application, the license number assigned to lottery winners, and the identity of parent firms, when relevant.

To obtain a set of product characteristics, we scraped all cannabis strains from the popular website `allbud.com`. We match strain names from that website to the strain names we observe in our data. We also use the Google Maps API to compute the distance between retailers and their suppliers, use the American Community Survey to obtain market-level data on demographic variables and obtain precipitation and temperature data at the county-month level from NOAA.

# 4 A Reduced-Form Test of Competing Bargaining Models

Before we proceed to outline the test we propose, based on the framework of section 2, we conduct a reduced-form test based on existing theoretical predictions and the variation in retail competition induced by the license lotteries discussed in section 3. We start from the following result, which summarizes conclusions obtained by Iozzi and Valletti (2014):

**Proposition 2** (Iozzi and Valletti (2014), Propositions 3 and 4)**.** *Suppose there is a single manufacturer, that retailers sell a single product, that the demand system of the Shubik and Levitan (1980) form, and retailers compete à la Bertrand. Then,*

- *If there is no repricing, the symmetric wholesale price may be decreasing, increasing, or non-monotonic in the number of retailers $R$, depending on the degree of product differentiation.*
- *If there is full repricing, the symmetric wholesale price is increasing in $R$.*

The relevance of this result for our purposes is that the no-repricing and full-repricing models make different predictions regarding the comparative statics of wholesale prices with respect to the number of retailers $R$.[6] Testing comparative statics related to the number of firms is notoriously challenging, as entry is a decision taken by firms. Fortunately, we can use the license lottery to obtain exogenous variation in the number of retailers and estimate the relationship between wholesale prices and downstream competition.

[6]Note that the partial-repricing regime is irrelevant in the context of this proposition, as the Iozzi and Valletti (2014) model focuses on single-product retailers.

To that end, we estimate the following specification

$$w_{jrt} = \sum_{k=1}^{4+} \beta_k \mathbf{1}\{N_{c(r)t} = k\} + \gamma_{jt} + \boldsymbol{x}'_{c(r)t}\boldsymbol{\theta} + \varepsilon_{jrt} \; . \tag{15}$$

In equation (15), $j$ indexes a product (i.e., a combination of a processor, a cannabis strain, and packaging), $r$ indexes a retailer, and $t$ indexes a year-month pair. The dependent variable $w_{jrt}$ is the wholesale price paid for product $j$ by retailer $r$ in month $t$. The explanatory variables of interest are the dummies $\mathbf{1}\{N_{c(r)t} = k\}$ which are equal to 1 if the number of retailers in retailer $r$'s city in month $t$ is equal to $k$. We include such dummies for 1, 2, 3 and 4 or more retailers (denoted $4+$). The coefficients $\gamma_{jt}$ are product-year-month fixed effects, and capture aggregate trends in the wholesale price of good $j$. Their inclusion in equation (15) implies that only cross-sectional variation within a product-year-month is used to estimate the $\beta_k$ coefficients. We also include controls $\boldsymbol{x}_{c(r)t}$.

We expect the dummies $\mathbf{1}\{N_{c(r)t} = k\}$ to be endogenous in equation (15). Unobserved demand and supply conditions likely affect both wholesale prices $w_{jrt}$ and the number of firms that chooses to enter the market. To identify the $\beta_k$ coefficients in (15), we leverage the license lottery. Following Borusyak and Hull (2023), we instrument $\mathbf{1}\{N_{c(r)t} = k\}$ with its expectation $\mathbb{P}(N_{c(r)t} = k)$, where this probability is computed by simulating counterfactual lottery outcomes.

Table 1 reports the results. Throughout all specifications, we find that wholesale prices are non-monotonic with respect to the number of retailers in a city. Specifically, the estimates in table 1 imply that wholesale prices are lowest under a retail monopoly, grow up to three retailers, and fall with four or more retailers. In light of Proposition 2, we interpret this result as evidence in favor of with no repricing.

The results of this section give us an indication of which bargaining model may be a better fit for our data. However, this test is rather limited. It is predicated on a theoretical result that makes several assumptions known to be or very likely to be false in our context: a single manufacturer, single-product retailers, and a linear demand system with a very restrictive form of product differentiation. Moreover, the test is based on a particular feature of our environment, namely the license lotteries. The test that we discuss next, based on the theoretical model from section 2, deals with all of these shortcomings. That framework allows for many multi-product manufacturers, multi-product

Table 1: Effect of Retail Competition on Wholesale Prices

| Dependent Variable: | Log(Price per gm) | | | |
|---|---|---|---|---|
| Model: | (1) | (2) | (3) | (4) |
| *Variables* | | | | |
| $N = 2$ | 0.0194*** | 0.0116 | 0.0158*** | 0.0159*** |
| | (0.0065) | (0.0081) | (0.0052) | (0.0058) |
| $N = 3$ | 0.0310*** | 0.0239** | 0.0330*** | 0.0471*** |
| | (0.0073) | (0.0105) | (0.0058) | (0.0089) |
| $N \geq 4$ | 0.0223*** | 0.0096 | 0.0220*** | 0.0245** |
| | (0.0080) | (0.0109) | (0.0052) | (0.0106) |
| Log(Median Income) | | 0.0059* | | 0.0158* |
| | | (0.0034) | | (0.0083) |
| Log(Adult Population) | | -0.0003 | | -0.0032 |
| | | (0.0021) | | (0.0045) |
| Share White | | -0.0053 | | -0.0205 |
| | | (0.0042) | | (0.0134) |
| Share Male | | 0.0336* | | 0.0261 |
| | | (0.0198) | | (0.0306) |
| *Fixed-effects* | | | | |
| Product-Seller-Month FEs | Yes | Yes | Yes | Yes |
| *Fit statistics* | | | | |
| Observations | 1,923,297 | 1,904,090 | 770,024 | 750,817 |

*Clustered (Year-Month-City) standard-errors in parentheses*
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

NOTES HERE.

retailers, and a general demand system. Moreover, we discuss sources of variation other than the license lottery that can be used to construct our tests. These features make the test of section 5 widely applicable.

## 5 Testing Competing Models

In this section we outline how we test competing bargaining models. Our tests build on and extend to a Nash-in-Nash model the tests developed by Backus et al. (2021) and Duarte et al. (2023).

We start from the tautology

$$c_{jt} = w_{jt} - \mu^*_{jt}$$

where $\mu^*_{jt}$ is the true upstream mark-up. Proposition 1 characterizes the mark-up function $\boldsymbol{\mu}_t(\boldsymbol{w}^*, \boldsymbol{\beta}; \mathcal{J}, \Omega)$ given a model of competition. It follows that under the true model

$$c_{jt} = w_{jt} - \boldsymbol{\mu}^*_{jt}(\boldsymbol{w}, \boldsymbol{\beta}; \mathcal{J}, \Omega) \tag{16}$$

We now introduce a set of exogenous variables $Z$. We partition these variables into $Z = (Z^{c'}, Z^{e'})'$, where $Z^c$ enters the marginal cost function and $Z^e$ is excluded from it. Let $c_{jt} = h(Z^c_{jt}) + \omega_{jt}$, where $\mathbb{E}[\omega_{jt} \mid Z] = 0$. This mean-independence assumption will be the basis for our tests.

Define

$$\omega^{\mu}_{jt} := w_{jt} - \boldsymbol{\mu}_{jt}(\boldsymbol{w}, Z^e, \boldsymbol{\beta}; \mathcal{J}, \Omega) - h(Z^c_{jt}) \ ,$$

where we have made the dependence of $\boldsymbol{\mu}_{jt}$ on $Z^e$ explicit. It follows from (16) that

$$\omega^{\mu^*}_{jt} = \omega_{jt} \ .$$

For any other model, however,

$$\omega^{\mu}_{jt} = \omega_{jt} + \Delta\mu(Z^e) \ ,$$

where $\Delta\mu := \mu^*_{jt} - \mu_{jt}$. Because $\Delta\mu$ is itself a function of $Z^e$ it follows that any model other than the true model will violate $\mathbb{E}[\omega^{\mu}_{jt} \mid Z] = 0$. This is the basis of the tests below.

We follow Backus et al. (2021) and Duarte et al. (2023) in adopting the Rivers and Vuong (2002) framework. Specifically, given a pair of models $\mu_1$ and $\mu_2$, we test the null hypothesis

$$H_0 : Q_1 = Q_2$$

where $Q_i := \mathbb{E}[\omega_{jt}^{\mu_i} f(Z)]$ for some (vector-valued) choice of instruments $f(Z)$. The test statistic is

$$T^{RV} = \frac{\sqrt{n}(\hat{Q}_1 - \hat{Q}_2)}{\hat{\sigma}_{RV}} \ ,$$

where $\hat{Q}_i := \hat{g}_\mu' W \hat{g}_\mu$, $\hat{g}_\mu := n^{-1} f(\boldsymbol{Z})(\boldsymbol{w} - \boldsymbol{\mu})$, and $W = n(f(\boldsymbol{Z})' f(\boldsymbol{Z}))^{-1}$. In these expressions, $f(\boldsymbol{Z})$ is a $n \times l$ matrix with the vectors $f(Z_i) \in \mathbb{R}^l$ in its rows, $\boldsymbol{w}$ is the vector of observed wholesale prices, and $\boldsymbol{\mu}$ is the vector of markups implied by model $\mu$.

# 6 Conclusion

To be written.

# References


Alviarez, V. I., Fioretti, M., Kikkawa, K., & Morlacco, M. (2023). *Two-sided market power in firm-to-firm trade* (Tech. Rep.). National Bureau of Economic Research.

Backus, M., Conlon, C., & Sinkinson, M. (2021). Common ownership and competition in the ready-to-eat cereal industry. *Working Paper*.

Bagwell, K., Staiger, R. W., & Yurukoglu, A. (2021). Quantitative analysis of multiparty tariff negotiations. *Econometrica*, *89*(4), 1595–1631.

Borusyak, K., & Hull, P. (2023). Nonrandom exposure to exogenous shocks. *Econometrica*, *91*(6), 2155–2185.

Crawford, G. S., Lee, R. S., Whinston, M. D., & Yurukoglu, A. (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica*, *86*(3), 891–954.

Crawford, G. S., & Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review*, *102*(2), 643–685.

De los Santos, B., O'Brien, D. P., & Wildenbeest, M. R. (2021). Agency pricing and bargaining: Evidence from the e-book market. *Kelley School of Business Research Paper*(18-90).

Draganska, M., Klapper, D., & Villas-Boas, S. B. (2010). A larger slice or a larger pie? an empirical investigation of bargaining power in the distribution channel. *Marketing Science*, *29*(1), 57–74.

Duarte, M., Magnolfi, L., Sølvsten, M., & Sullivan, C. (2023). Testing firm conduct. *arXiv preprint arXiv:2301.06720*.

Ellickson, P. B., Lovett, M. J., Sunada, T., & Kong, P. (2022). Private labels and retailer profitability: Bilateral bargaining in the grocery channel. *Available at SSRN 3045372*.

Gowrisankaran, G., Nevo, A., & Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, *105*(1), 172–203.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, *103*(1), 145–177.

Ho, K., & Lee, R. S. (2017). Insurer competition in health care markets. *Econometrica*, *85*(2), 379–417.

Hollenbeck, B., & Uetake, K. (2021). Taxation and market power in the legal marijuana industry. *The RAND Journal of Economics*, *52*(3), 559–595.

Horn, H., & Wolinsky, A. (1988). Bilateral monopolies and incentives for merger. *The RAND Journal of Economics*, 408–419.

Iozzi, A., & Valletti, T. (2014). Vertical bargaining and countervailing power. *American Economic Journal: Microeconomics*, *6*(3), 106–135.

Miravete, E. J., Seim, K., & Thurk, J. (2018). Market power and the laffer curve. *Econometrica*, *86*(5), 1651–1687.

Rivers, D., & Vuong, Q. (2002). Model selection tests for nonlinear dynamic models. *The Econometrics Journal*, *5*(1), 1–39.

Shubik, M., & Levitan, R. (1980). *Market structure and behavior*. Harvard University Press.