How big is the "lemons" problem? Historical evidence from French appellation wines

Pierre Mérel[†], Ariel Ortiz-Bobea[‡], and Emmanuel Paroissien[⋆]

[†] Agricultural and Resource Economics, University of California, Davis, Davis, CA. Email: merel@primal.ucdavis.edu

[‡] Charles H Dyson School of Applied Economics and Management, Cornell University, Ithaca, NY. Email: ao332@cornell.edu

[⋆] INRA, Rennes, France. Email: Emmanuel.Paroissien@inra.fr

This is a preliminary version. Please do not cite.

Abstract

We provide an empirical measure of the economic surplus loss arising from the failure of a competitive market to supply quality in the presence of asymmetric information. When consumers cannot observe product characteristics at the time of purchase, incentives for atomistic producers to supply costly quality may be suppressed. We use natural variation in wine prices across administrative districts around the enactment of pioneering regulations aimed at resolving asymmetric information problems in the French wine market to identify related welfare losses. Results from a panel analysis indicate large potential losses from the quality-related market failure, suggesting an important role for credible certification schemes.

JEL codes: D82, N54, Q18
Keywords: asymmetric information, quality uncertainty, wine, France, appellation

In his seminal paper "The market for lemons: quality uncertainty and the market mechanism," George Akerlof formalized the notion that a consumer's inability to ascertain objective quality differences in products may "drive the good product out of the market," resulting in a socially undesirable outcome (Akerlof, 1970). Akerlof's insight was that if buyers cannot distinguish good products from bad, they will value a product as having average quality. If sellers of the good product have reservation prices that, despite being lower than buyers' valuation of it, lie above buyers' valuation of the average-quality product, they cannot profitably trade with them. In equilibrium, the bad product is sold yet the good product remains in the hands of sellers, despite having higher social value in those of potential buyers. That is, asymmetric information can suppress mutually beneficial trade.

At the time of its publication, Akerlof's piece was famously dismissed by some economists as either trivial or wrong. Half a century later, no economist would argue that Akerlof's description of the quality-related market failure was conceptually flawed. In the meantime, a growing body of empirical work has sought evidence of the lemons effect in various real-world contexts, sometimes successfully. Yet, evidence on the size and welfare significance of the effect is hard to come by. One notable exception is insurance markets, which have been the object of much empirical scrutiny in recent times, but even there the available evidence on social welfare effects has been limited (Bundorf et al., 2012). Estimates of the welfare effects of

selection in insurance markets, when available, typically fall within a few percent of market value (Handel, 2013).

The present paper relies on a different empirical setting—the French wine market—to identify the welfare impacts of the lemons effect. Using observational price data and a policy change meant to mitigate informational asymmetries, we provide evidence of a large welfare effect channeled through adverse selection in quality provision by atomistic sellers.

Let us start by highlighting the fact that, in addition to its clear potential to suppress mutually beneficial trade—illustrated in Akerlof's paper in the context of an exchange economy—asymmetric information about product quality can also deter *production*. Intuitively, if buyers cannot tell quality differences at the time of purchase, and quality is costly to supply, atomistic producers have no incentive to supply it. The resulting market equilibrium may involve an exchange of goods, but there are gains forgone from not producing—and trading—higher-quality goods instead. Like Kim (1985), we view the deterrence to the supplying of valuable quality as another manifestation of Akerlof's lemons effect.

We argue that such a lack of incentives to supply quality was at play in the French wine market during the decades preceding the adoption of a 1935 law that codified production rules and implemented official controls for fine wines claiming a reputable geographical appellation—like *bordeaux* or *bourgogne*. We show that this pioneering law, the first of its kind to be adopted in the world and the enduring template for any regulation pertaining to geographical indications, profoundly and durably changed the nature of the French wine market. Our analysis, which involves a careful counterfactual comparison of district-level wine prices before and after the reform, also reveals the extent of the market failure preceding its adoption.

The wine market is an ideal setting to study the effects of quality-related market inefficiencies. Wine is a highly differentiated product, with the area of origin potentially playing a salient role in signaling quality. Yet it is difficult for the average consumer to know quality at the time of purchase, even with a geographical indication. To the extent that the use of the appellation is free, incentives to free-ride on a region's longstanding reputation are large as production is atomistic and costs vary greatly according to the varietal planted, the type of terrain, or the techniques used to turn grapes into wine. Hence, wine producers may be tempted to plant high-yielding but low-quality grapes on inappropriate terrains and resort to subpar wine-making techniques while claiming a theretofore reputable

origin. And indeed the history of wine production—France is no exception—is riddled with anecdotes of such deceptive but profitable behavior. Whether these well-documented anecdotes add up to economically meaningful effects, and if so, whether some form of government intervention may be effective at correcting them, is a more debatable proposition, which the present paper seeks to address.

To identify the extent of the quality-related market failure, we assemble a panel of yearly average wine prices received by producers in each department—a French administrative unit roughly the size of a US county—for the period 1907–1969. We regress the departmental price of wine on the share of a department's vineyards eligible for recognition under *appellation d'origine contrôlée* (hereafter AOC), the official designation for appellation wines created by the 1935 reform.[1] We control for time-invariant unobservables through locational fixed effects and for yearly shocks through spatially-differentiated time effects. Because it took time for the administration to define the 213 AOCs present during our sample period, our measure of AOC eligibility does not go from zero to its final value within a year, but instead grows as more AOCs are recognized over time. The fact that departments have varying shares of vineyards eligible for an AOC (many having shares equal to zero, others one, and many others in between) and the temporal roll-out of the reform allow us to flexibly control for potentially confounding factors through year fixed effects differentiated by broad wine region. We also control for wine production to capture swings in wine prices arising from weather shocks and the possibility that the reform may have reduced wine output.

Our regression analysis yields an estimate of the rate of increase of the average departmental wine price with respect to the share of vineyards eligible for AOC equal to 45%. Not all eligible vineyards claimed an AOC however, and our estimate implies that actual AOC recognition led to an increase in the price of appellation wines roughly equal to the size of the average wine price. This dramatic appreciation suggests that appellation wines for which production was ultimately regulated had been produced at an inefficiently low quality prior to regulation, consistent with historical accounts of widespread abuse in the appellation wine market in the

[1]There were several legislative attempts to define appellation wines prior to 1935. None of them included official controls or a systematic definition of production requirements. In many cases, definitions merely included broad geographical delimitations, which encouraged free-riding on other important aspects of quality provision within the delimitated zones, and led to a worsening, not an improvement, of the asymmetric information problem (Capus, 1947). This may partly explain why Haeck et al. (2018), who study the impact of these pre-1935 reforms, obtain mixed results.

decades leading to the reform. Importantly, we do not find any evidence that the reform decreased wine production, which implies that the price increase cannot be attributed to a reduction in the quantity of wine sold. We are also able to reject a competing hypothesis according to which the price increase was the result of the *déclassement* of wines, that is, the denial of an appellation label for wines sold under an appellation prior to the reform.

We provide three separate pieces of evidence in support of the "parallel trends" assumption necessary for a causal interpretation of our estimates. First, we show that eventual AOC recognition had absolutely no effect on price patterns during the pre-reform period going from 1907 to 1936. Second, reducing our sample to the period 1938–1969, during which the majority of AOCs had already been established, we show that the residual variation in AOC recognition leads to estimates that are comparable to those from the full sample of years. Unobserved factors would thus need to be systematically correlated with the timing of AOC recognitions for the estimated effects to be spurious. Third, we remove from our sample departments with low eventual AOC share, that is, we use the intensive margin of AOC propensity as the source of identifying variation. Although the resulting sample is less than half the size of the original one, the estimate is extremely close.

Importantly, our main estimated effect is also robust to the removal from our sample of years around WWII when the French wine market is believed to have been disrupted, notably due to forced government procurement below market prices.

At the end of our study period, the share of French vineyards eligible for AOC recognition reached 30.5%. Together with our estimated effect on the average wine price, this figure implies a welfare loss of about 14% in the French wine market—inclusive of all wines—due to asymmetric information prior to the reform. This value represents a *gross* welfare loss in the sense that it does not account for the added cost of quality-enhancing practices required for wines sold under an AOC label. While these cost increases could be substantial, the fact that a large share of eligible producers decided to claim an AOC—as opposed to producing cheaper wines—suggests that the policy was beneficial to wine producers, and welfare-enhancing. Irrespective of the actual costs of supplying quality, our gross welfare measure reveals the considerable size of the latent unsatisfied demand for higher-quality wine under asymmetric information, suggesting that the lemons problem may severely affect market performance in vertically differentiated markets with atomistic supply, even when the volume of trade is large.

This paper directly relates to a rich empirical literature seeking evidence of adverse selection in real-world markets. In line with Akerlof's original setting, some of this literature has focused on the trade of durable goods such as used vehicles (Bond, 1982; Genesove, 1993; Lewis, 2011). None of these papers finds strong evidence of adverse selection, suggesting a limited role of asymmetric information. A richer strand of literature has investigated insurance markets (Puelz and Snow, 1994; Cutler and Reber, 1998; Cawley and Philipson, 1999; Chiappori and Salanié, 2000; Cardon and Hendel, 2001; Finkelstein and Poterba, 2004; Einav et al., 2010; Einav and Finkelstein, 2011; Handel, 2013; Hackmann et al., 2015; Panhans, 2019). In these markets, sellers (insurance firms) have less information than buyers of insurance because they cannot fully observe buyers' riskiness. There, the evidence on adverse selection is mixed (Einav and Finkelstein, 2011), but welfare effects of adverse selection, when present, typically fall within a couple of percents of market value (Handel, 2013; Hackmann et al., 2015). Closer to our setting, Jin and Leslie (2003) examine the effects of quality information provision on firms' choices of quality in the context of restaurant hygiene. Like us, they exploit a policy change that mitigated information asymmetry and find evidence of quality improvements, although they are unable to make welfare claims.

Our paper also contributes to a broader literature on the impact of information disclosure on economic outcomes. Experiments have shown how improved access to and control of information can increase market efficiency by lowering search costs and limiting corruption (Jensen, 2007; Jensen and Miller, 2018; Andrabi et al., 2017; Duflo et al., 2013), or instead generate perverse selection effects (Dranove et al., 2003). Finally, our paper relates to a literature that has exploited wine quality signals to analyze the functioning of quality-differentiated markets, their trade (Crozet et al., 2012), and the formation of prices (Ashenfelter, 2008; Ali et al., 2008).

The rest of the paper is organized as follows. Section 1 provides some historical institutional background. In Section 2, we formalize the asymmetric information problem in the context of endogenous quality provision using a simple model of vertical differentiation. (Alternative models are presented in Appendices.) Importantly, we highlight how the model can be brought to the data and key parameters estimated to derive meaningful gross welfare effects using available average wine price data. Section 3 exposes our identification strategy, the construction of our dataset, and our empirical results, including tests of the parallel trends assumption and a series of robustness checks. Section 4 concludes.

1 Historical and institutional background

The AOC system was created by a 1935 French law as the outcome of a longstanding debate on the recognition and preservation of premium quality wine-producing areas, known as appellations. Two issues were particularly debated: (i) the geographic borders of these areas and (ii) the set of eligible vineyard and wine-making practices. The search for a consensus on these questions caused a series of regulatory trials and errors throughout the 20th century, leading to the coexistence of a set of certifications of origin, AOC lying at the top of the hierarchy.

Before any regulation on wine appellations was adopted, France's most renown vineyards (*vignobles*, meant here as potentially large sets of parcels), whose place names were already used to identify the wines produced therein, suffered from free-riding and malpractice. These problems became widespread during the acute production shortage of the late 19th century.[2] This crisis generated strong incentives among wine producers to increase production while lowering quality.[3] Malpractice was so prevalent that in 1889, French authorities had to pass a law defining wine as the exclusive produce of grape juice fermentation. During this episode, quality vineyards were especially harmed since the general trend was to produce lower quality wines at higher yields. Furthermore, at the time there existed no legal definition of appellation wines. Unsurprisingly, counterfeiting was common as famous names were often usurped by producers located in other wine regions, or were used without consideration for the production techniques and attendant wine characteristics that had brought reputation to the place.

In 1905, France adopted its first general law on the prevention of fraud and falsification. Although its scope was much broader than the protection of wine appellations, the law provided a mechanism by which the French administration would take on the task of delineating the geographical limits of each wine appellation.[4] Those boundaries were to be defined by administrative decrees. A few appellations were thus delimited, starting with the *champagne* appellation in 1908, followed by *banyuls*, *cognac*, and *armagnac*. The administration then delimited *clairette de Die* in 1910 and *bordeaux* in 1911 (Humbert, 2011). Unfortunately, this

[2] In the 1860s, a pest imported from America called phylloxera started to ravage French vineyards, eventually causing production to be cut by half between 1875 and 1890.

[3] A common way to increase volume while maintaining the alcohol content of wine was to add sugar to the must and dilute wine with water. Another way was to fabricate wine from raisins.

[4] This task was defined in a 1908 amendment to the 1905 law.

top-down definition of appellation regions proved unsatisfactory to many stakeholders. It is often cited as a leading cause of the Champagne Riots of 1911, as producers in excluded regions felt they had been wrongly denied the appellation. Administrative delineations also failed in the Bordeaux region. In addition to generating political unrest, administrative delineations had a fundamental weakness: they established a legal right to utilize a place name based solely on delimitations at the level of the municipality, irrespective of the type of terrain, grape varietal, or production practices. Not surprisingly then, unscrupulous producers located in eligible regions started to market mediocre wines under famous appellations. This situation raised concerns among higher-quality producers who were often supportive of precise eligibility conditions for appellation wines (Capus, 1947).

A 1919 law removed the authority to define appellation wines from the French administration and gave it to the courts. Any stakeholder who thought they were being hurt by the abusive use of a place name could file a lawsuit. Courts were given the right to not only define geographical boundaries but also to take account of "local, loyal, and constant uses." Unfortunately, most judges refrained from defining production practices, and in effect, for most appellations the court only specified geographical boundaries, just as the former administrative decrees.[5] As a result, in the early 1930s most appellations only had requirements pertaining to the eligible area. This period also saw a rise in the number of new appellations claimed by producers as a way to escape the stringent production controls applicable to ordinary wines starting in 1931 with the *Statut Viticole*. This situation led to what is known as the "appellation scandal," that is, the proliferation of unwarranted appellations, which further eroded the reputation of historical appellations.

The 1935 law introduces a new category of so-called "controlled origin appellations" (*appellations d'origine contrôlée*, or AOC), without—at first—eliminating existing appellations. These new appellations are to be defined by administrative decrees. But unlike the early administrative delimitations, the provisions of the AOC decree are not dictated by the administration. Instead, the decree sanctions a set of production requirements, including detailed geographical boundaries at the parcel level, that emanate from a committee composed, by order of importance, of representatives of local wine associations and wholesalers, members of Parlia-

[5]Another law passed in 1927 explicitly allowed courts to include a list of specific grape varieties as well as soil restrictions in the definition an appellation, but these precisions were left optional, and very few judgements included restrictions on terrain or varietal (Ministère de l'agriculture, 1937; Capus, 1947).

ment, and representatives of the administration—the CNAO, *Comité national des appellations d'origine*. As such, the definition of the requirements applicable to each AOC is left to a technical body of experts that includes representatives of each wine region.[6] In contrast to existing appellations, now referred to as "plain appellations" (*appellations simples*), AOCs are subject to official control. Wines can claim an AOC if they are grown on an eligible parcel, according to the specified practices, and meet a set of criteria pertaining to, e.g., alcohol content. The AOC is not compulsory in the sense that producers may elect to sell their wines as ordinary wines, or under a plain appellation (without control) if they can claim one based on location. Typical requirements for an AOC, beyond geographical area and terrain, are the grape varietal, the specification of a maximum yield per hectare, and minimum levels for alcohol and sugar contents.[7]

Soon after the 1935 law, many appellations were officially recognized by an AOC decree: 77 AOCs were created in 1936 and 58 others in 1937. These new AOCs did not exactly replace the former appellations of the same names: both an AOC and a plain appellation could coexist under the same name in the same region. For instance, after the creation of the *bordeaux* AOC in 1936, Bordeaux wines that did not meet the strict requirements of the AOC could still be sold under the plain appellation. This coexistence of both plain and controlled appellations, known as the "double appellation regime," although arguably confusing, was necessary to garner political support for the new system as it allowed producers willing to claim an AOC to transition to the new requirements. However, this regime was soon to be abolished.

A first law passed in 1938 allowed the CNAO to forbid the use of a plain appellation at the request of the most representative local producer organization. This option was immediately adopted in many small, upper-quality regions, and by the end of 1939, half of the AOCs had gotten rid of their plain appellation counterparts. However, large regional appellations like *bordeaux* and *bourgogne* survived the creation of their AOC counterpart as no consensus was found in their respective local unions in favor of abolition. This situation was put an end in

[6]The CNAO was initially financed by a tax on the sales of AOC wines of 2 francs per hectoliter. Its agents were sent to delimit each AOC at the parcel level and to control production conditions.

[7]In the late 1920s, some appellation wines were produced at very high yields, between 120 and 200 hectoliters per hectare but with only 7% of alcohol content in volume (Capus, 1947). The minimum alcohol content for AOC wines was typically set to between 10% and 15%, and the maximum yield between 20 and 50 hectoliters per hectare. These figures are still current standards for AOC wines.

1942 when a new law granted the CNAO the right to unilaterally suppress a plain appellation wherever an AOC also existed under the same name. All remaining duplicate appellations were eliminated the next year. Thus, the only surviving plain appellations were those for which no AOC had been created. The AOC label quickly became the standard for premium quality wines and by 1940 the production of AOC wines already exceeded that of plain appellation wines (Humbert, 2011).

By 1940, 151 different AOCs had already been created, a testimony to the large amount of regulatory work undertaken by the CNAO.[8] Nonetheless, the CNAO was led to reject several AOC requests, as some less-known vineyards were found too heterogeneous and therefore unfit to bear the AOC label.[9]

Note that AOC delineations can be cumulative as the structure of the AOC system is inherently hierarchical. That is, a given parcel may be eligible for several appellations. For instance, a parcel located on appropriate terrain in the municipality named Pauillac would be eligible for the following appellations, ranked from the most common to the most exclusive: *bordeaux*, *bordeaux supérieur*, *haut-médoc*, and *pauillac*.

At the end of 1949, a new intermediary quality label called VDQS (*vin délimité de qualité supérieure*) was introduced to reward the best non-AOC vineyards. Three different certifications of origin then coexisted for a few years: plain appellation, VDQS, and AOC. The first two remained smaller in volume, whereas the AOC label established itself as the standard certification for premium quality wine.[10] From the years following the 1935 law to the year 1969 that marks the end of our observation period, AOC wines represented on average between 10 and 15% of total French wine production.

[8] The 300^{th} wine AOC was created in 2012. The concept of AOC has been extended in 1990 to all agricultural products such as cheese, fruits, or olive oil, and is now in use throughout the European Union.

[9] The examination of an application included a tasting session and an assessment of the reputation of the wines produced in the candidate region (Humbert, 2011).

[10] In the 1950s and throughout the 1960s, the production of VDQS wines only represented between one third and one half of the production of AOC wines (Humbert, 2011). The production of plain appellation wines also remained about half that of AOC wines. The plain appellation and VDQS labels were abandoned in 1973 and 2011, respectively.

2 A model of the wine market with endogenous quality

We model wine production at the level of a French department. Vineyard acreage is assumed to be inelastic, and we further assume that yields are fixed (they may vary across space, but are not affected by regulation). As we show in the empirical section, these assumptions, though perhaps unexpected, appear warranted by our data. Since there are no quantity effects, we can focus on the impact of regulation on wine quality.

For simplicity, we assume that there are two broad categories of wines, (i) ordinary wines grown in places where climate and soils can only yield mediocre wine, and (ii) appellation wines grown in places endowed with beneficial natural factors such as climate and soils, the effects of which may be further enhanced by appropriate production practices, such as varietal choice, winemaking techniques, etc. The second category of wine is distinguished from the first at wholesale and retail by the prominent use of the name of the place from which the wine originates—the appellation. In a department, there may be more than one appellation, but we assume that consumers' valuation of appellation wine is uniform across appellations before the reform. In contrast to appellation wines, ordinary wines are assumed to have a fixed quality that cannot be enhanced through costly practices.[11]

We further assume that there are many identical consumers, each with unit demand for wine, and that there are more consumers than units of wine produced.[12] Therefore, wines are sold at a price equal to their consumer valuation, and some consumers are not served. The consumer valuation of ordinary wines is denoted p_0, and that of appellation wines, when no costly production practices are used, is denoted p_1.

Note that before any regulation on production practices is enacted, a market equilibrium cannot involve any costly practices for appellation wines. The reason is that a single producer engaging in such practices would have an incentive to shirk since consumers cannot tell quality differences among appellation wines at the time of purchase, and there are many wines claiming the same appellation.[13]

[11] Technically, we could allow for the possibility of quality enhancement, but the free-rider problem would prevent any producer from profitably pursuing it.

[12] This assumption may seem at odds with the observation that in some years, there may exist production surpluses, leading to very low wine prices. Our model is to be understood as a static representation of a multi-year market equilibrium where production is inelastic and weather shocks average out.

[13] One implicit assumption is that individual producers of appellation wines cannot reliably signal

We assume that $p_1 \geq p_0$, that is, appellation wines cannot be of lower quality than ordinary wines.

We denote by s_1 the share of appellation wine produced and by $s_0 = 1 - s_1$ the share of ordinary wine produced. Although appellation and ordinary wines may be sold at different prices since appellation wines are distinguishable by their place name, in the data we only observe the average price of wine, $p_m \equiv p_0 s_0 + p_1 s_1 = p_0 + s_1(p_1 - p_0)$.

After the reform, the use of a place name is restricted, for wines bearing the AOC label, to wines produced according to certain quality-enhancing practices. For plain appellation wines no specific production techniques are mandated. The reform therefore creates a difference between two types of appellations, plain appellations and AOCs, that may sell at different prices.

The reform leaves consumers' valuations of ordinary wines and plain appellations unaffected. In contrast, wines sold under the AOC label, which were previously sold as plain appellations, have a (weakly) higher valuation after the reform, say $p_2 \geq p_1$. Denote by s_2 the share of wine eligible for AOC after regulation. We assume $s_2 \leq s_1$, with the strict inequality corresponding to the case where not all wine previously sold under appellation can claim an AOC.[14]

After regulation, we can thus write the average price of wine in a department as:

$$\begin{aligned} p_m &= (1 - s_1)p_0 + (s_1 - s_2)p_1 + s_2 p_2 \\ &= \underbrace{p_0 + s_1(p_1 - p_0)}_{(A)} + \underbrace{s_2(p_2 - p_1)}_{(B)}. \end{aligned} \tag{1}$$

The terms (A) in Equation (1) depend only on a department's appellation share and exogenous characteristics, but not on regulation, while term (B) depends on the extent of regulation. The effect of the reform on the department's wine price is $\Delta p_m \equiv s_2(p_2 - p_1)$.

quality to consumers, perhaps because of the very large number of producers in a given appellation region, which makes it very difficult for a single producer to create a reputation beyond the collective reputation of the appellation. For instance, in 1935 there were more than 65,000 wine producers in Gironde, the department where *bordeaux* is produced.

[14]We could further differentiate the valuations of plain appellations and AOC wines before the reform, based on the idea that wines declared eligible for an AOC likely benefit from different natural factors than those only worthy of a plain appellation. This refinement would complicate the model without adding anything to our argument or the interpretation of our regression coefficients.

Figure 1 Welfare effects of asymmetric information in the appellation wine market

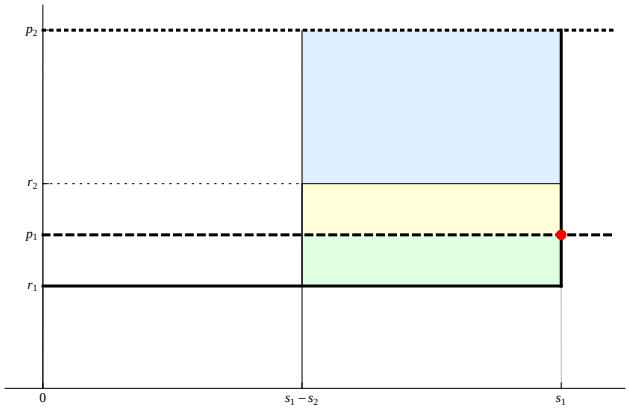

Fundamentally, we are interested in an empirical measure of the value $p_2 - p_1$, which captures consumers' valuation of the quality of an appellation wine that fails to be supplied under asymmetric information. If there is no lemons effect, i.e., quality does not improve after the reform, then $p_2 = p_1$. In addition, the product of the price increase $p_2 - p_1$ by the quantity of AOC wine directly translates into a partial (or gross) welfare increase:

$$\Delta \text{GW} = Q s_2 (p_2 - p_1) \tag{2}$$

where Q denotes total wine output. Note that in our model with perfectly elastic demand for wines of a given quality and perfectly inelastic supply, all welfare accrues to producers. However, our measure of welfare improvement is partial because it does not account for the cost of quality-enhancing practices adopted on the share s_2 of production.

Figure 1 depicts the gross and net welfare losses from asymmetric information, in the case where $s_2 = \frac{s_1}{2}$, that is, only half of appellation wine production is deemed worthy of an AOC. Total wine output is normalized to one. Since the price of ordinary wines does not change with regulation, only the market for appellation wine is depicted. The average cost of supplying appellation wine is assumed to be constant and equal to r_1 while that of supplying AOC wine is assumed to be constant and equal to $r_2 > r_1$. The net welfare loss from asymmetric information, which is resolved by regulation, is the difference between the area shaded in blue (which represents the welfare from the trade of regulated wine under full information) and the area shaded in green (the welfare from the trade of this wine under

asymmetric information). The gross welfare loss only relates to differences in consumer valuations (or market prices) and is given by the sum of the areas shaded in blue and yellow. The red dot depicts the equilibrium price of appellation wine under asymmetric information.

Importantly, Equation (2) can also be used to derive the relative change in gross welfare

$$\frac{\Delta \text{GW}}{\text{GW}} = \frac{Qs_2(p_2 - p_1)}{Q\left[(1-s_1)p_0 + s_1 p_1\right]} = \frac{\Delta p_m}{p_m} \approx \Delta \log p_m \tag{3}$$

where $\Delta \log p_m$ represents the change in the department's log average price attributable to regulation. Thus, a regression of $\log p_m$ on the share s_2 of a department's wine production eligible for a controlled appellation (with appropriate covariates to control for confounding factors) will yield the partial derivative $\frac{\partial \log p_m}{\partial s_2}$, which multiplied by the ultimate share of production eligible after the reform becomes a predictor of $\Delta \log p_m$ and thus of $\frac{\Delta \text{GW}}{\text{GW}}$. We can further interpret the coefficient on s_2, say β, as the price premium relative to the average price of wine. This is because $\log p_m = \log\left(p_0 + s_1(p_1 - p_0) + s_2(p_2 - p_1)\right)$, and thus $\beta \equiv \frac{\partial \log p_m}{\partial s_2} = \frac{p_2 - p_1}{p_m}$. Given Equation (3), the coefficient β itself has a clear welfare interpretation: it is the relative rate of increase of gross welfare with respect to the share of wine production eligible for quality certification.

Before moving to the empirical part of this study, which is concerned with obtaining an unbiased estimate of β, we wish to make four remarks. First, the mere observation that the price of eligible appellation wines rose after the reform—assuming we could observe appellation wine prices, which we do not—would not be sufficient to conclude that the reform had had any effect on wine quality or welfare. Indeed, such a finding could be the result of the market moving from a pooling equilibrium where all appellation wines are sold at the same average valuation to a separating equilibrium where higher-quality wines distinguished by the AOC label command a higher price while lower-quality wines sell at a lower price. Although such an equilibrium shift would have obvious distributional impacts, efficiency would not be affected as long as qualities were exogenously determined and unaffected by the reform. It is thus important to estimate the effect of the reform on *average* wine prices—which we do observe—to test the welfare-enhancing character of the reform. Intuitively, the fact that the average wine price is found to increase with the share of vineyards eligible for AOC recognition suffices to establish the welfare-creating effect of the reform—at least ignoring increased

production costs.[15]

Second, not all wines eligible for AOC recognition are marketed as AOC wines. In particular, some producers eligible based on vineyard location may supply baseline quality valued at p_1 because the associated costs would make AOC production unprofitable for them. The coefficient β should then be interpreted as capturing the average valuation difference for eligible wines (accounting for the fact that some of them remain plain appellations) i.e., an intention-to-treat effect. Formally, denote by $0 \leq \kappa \leq 1$ the share of eligible wine actually sold under AOC. Then, $p_m = (1 - s_1)p_0 + (s_1 - s_2\kappa)p_1 + s_2\kappa p_2 = p_0 + s_1(p_1 - p_0) + s_2\kappa(p_2 - p_1)$, $\Delta p_m = s_2\kappa(p_2 - p_1)$, and $\Delta GW = Qs_2\kappa(p_2 - p_1)$. Therefore, the coefficient on the eligible share, β, can still be used for welfare inference.[16]

Third, the derivation of ΔGW in Equation (2) assumed that all consumers have identical tastes. In Appendix A, we formally derive the expected welfare effects from wine regulation in a model where consumers have different tastes for quality. Importantly, we show that the gross welfare measure ΔGW derived above constitutes a lower bound to the gross welfare change when consumers are heterogenous in their valuation of quality. The intuition behind this result is that the valuation of the marginal consumer of AOC wine is lower than that of inframarginal consumers, and prices reflect marginal valuations.

Fourth, although the model in Equation (1) uses the share of wine *production* eligible for AOC (s_2) as a determinant of the average price of wine p_m, in our empirical implementation we use the share of *acreage* in vineyards eligible for AOC rather than a volume share. To the extent that yields are not affected by the reform (we provide empirical evidence in support of this fact), using the acreage share in

[15] If consumers are heterogenous with respect to their taste for quality, then, as we show formally in Appendix A.2 the average price of wine will rise without any quality changes if some wine previously sold under an appellation becomes ineligible and is sold as ordinary wine (there is historical evidence of such *déclassement* at least in the Bordeaux region). We show that in that case, welfare would also increase as quality-valuing consumers are able to select into consumption of a higher-quality appellation wine. Therefore, the relative change in average wine price has the same sign as the relative change in gross welfare—though not necessarily the same magnitude. Empirically however, we are able to reject *déclassement* as a driver of the increase in the price of wine thanks to the rollout of the reform, which temporarily allowed producers to continue using a place name even if they did not meet the production requirements set forth in the relevant AOC decree. See Section 3.3.3.

[16] If, in addition, a share $1 - \kappa$ of wines eligible for AOC recognition end up being sold as ordinary wines rather than plain appellations (perhaps because there is no plain appellation after the reform) the average valuation for ordinary wine will increase to $\tilde{p}_0 = \frac{(1-s_1)p_0+s_2(1-\kappa)p_1}{1-s_1+s_2(1-\kappa)}$, so that the average wine price will still be $p_m = (1-s_1)p_0+s_2(1-\kappa)p_1+(s_1-s_2)p_1+s_2\kappa p_2 = p_0+s_1(p_1-p_0)+s_2\kappa(p_2-p_1)$. This case is functionally similar to the previous one.

place of s_2 affects the structural interpretation of our regression coefficient if yields differ for ordinary and appellation wines. Denoting σ_1 the share of vineyard acreage initially under appellation, σ_2 the share of vineyard acreage eligible for AOC, y_0 the yield of ordinary wine and y_1 the yield of appellation wines (assumed to be unaffected by the reform), the change in the average wine price can be written as $\Delta p_m = s_2(p_2 - p_1) = \sigma_2 \frac{(p_2-p_1)y_1}{(1-\sigma_1)y_0+\sigma_1 y_1}$. To the extent that $y_1 \leq y_0$, the multiplier on the acreage share is therefore interpretable as the valuation increase modulated by the ratio of the appellation yield to the average yield. In the context of a regression of $\log p_m$ on σ_2, the coefficient of interest is interpretable as the relative rate of increase of gross welfare with respect to the share of vineyard acreage eligible for AOC.

3 Empirical analysis

3.1 Data

Our dataset combines several sources. Departmental average wine prices, areas in vineyards, and wine production come from France's *Statistique agricole annuelle*, a yearly publication of the Ministry of Agriculture available in print for the historical period. We focus on the period 1907-1969. This window excludes the period when phylloxera destroyed most of France's vineyards starting in the late 1850s. It further excludes an ensuing period of generalized fraud through wine adulteration, which ended with the adoption of the 1905 Law against fraud and falsification and the creation of the fraud repression service in 1907. We end our analysis in 1969, one year before the adoption of the first European regulation pertaining to the common organisation of the market in wine (Council of the European Communities, 1970).

We rely on several other sources of information to construct our main regressor, namely the share of vineyards in a department eligible for AOC at a given point in time. The first one is the set of governmental decrees defining each AOC pursuant to the 1935 law. These decrees provide information on the administrative area eligible for an appellation, typically by stating which municipalities (*communes*) are eligible for a given appellation (this area may cross departmental boundaries).

Historical records of which parcels within an eligible *commune* are eligible for an AOC are kept in the cadastral archives of each municipality. Reconstructing the historical record of eligible parcels would require visiting each municipality, which is prohibitive. Instead, we make use of a recent effort by France's *Institut*

national de l'origine et de la qualité (INAO) to map out eligible parcels using GIS tools. INAO has released a series of shape files covering a large share of France's current AOCs (notable exceptions include *champagne* and *vins doux*).[17] Among these, we first select AOCs that existed during the period of investigation (i.e., we exclude newer AOCs). For those AOCs that existed but are not part of the INAO data, by default we select the entire surface of the municipalities listed in the historical decrees of AOC recognition. Finally, because eligible parcels often include land not actually in vineyards (for instance, they may include hedgerows or access roads, or, for AOCs not covered by the INAO files, the entire municipality), we cross these delimitations with a land use raster file that shows pixels planted in vineyards in the years 1990, 2000, 2006, or 2012. The land use information comes from satellite imagery and these are the only years for which it is available. We cross the two files by first rasterizing the INAO shape file and then overlaying it over the land use file.

For AOCs that are covered by the INAO shape files, the resulting raster file therefore indicates the pixels eligible for the AOC as of 2016 while having been grown in vineyards in at least one of the four years for which we have explicit land use data. For the few appellations not covered by the INAO file, the resulting file indicates all pixels within an eligible municipality that are planted in vineyards in at least one of the years 1990, 2000, 2006, or 2012. The areas covered by these pixels are then summed up across departments' administrative boundaries. (Each pixel covers 1 ha of land.)

Finally, to construct the AOC eligible share at the level of the French department, we divide the area covered by pixels eligible for at least one AOC (while being grown in vineyards) in a department by the maximum of the area planted in vineyards during the period 1907–1969, which comes directly from the historical record in the *Statistique agricole annuelle*. This calculated share represent our best estimate of the true historical share of vineyards eligible. Similarly, we construct the departmental share of vineyards eligible for, say, five or more AOCs, by only selecting pixels that appear in five or more AOC delineations. For each AOC, we use the year following the year of enactment of the decree as the starting date for counting AOC eligibility. Since the first round of decrees were enacted in 1936, this rule implies that our regressor of interest may start taking non-zero values in the year 1937.

More formally, denote by i a department, by t a year, by l an AOC, and by p a

[17] Based on personal communication by one of the authors, it is not clear when the complete set of AOCs will be made available by INAO.

one-hectare pixel. Let us further denote:

$$\mathbb{1}_t^l = \begin{cases} 1 & \text{if AOC } l \text{ is recognized as of year } t \\ 0 & \text{otherwise} \end{cases};$$

$$\mathbb{1}_p^l = \begin{cases} 1 & \text{if pixel } p \text{ belongs to AOC } l \\ 0 & \text{otherwise} \end{cases};$$

$$\mathbb{1}_p = \begin{cases} 1 & \text{if pixel } p \text{ was grown in vineyards in 1990, 2000, 2006, or 2012} \\ 0 & \text{otherwise} \end{cases}.$$

Given that we start counting recognition in the year following the AOC decree, the indicator $\mathbb{1}_t^l$ equals zero from 1907 until the year in which the decree for AOC l is enacted, and one thereafter.

Using this notation, we define $\Delta_p^l \equiv \mathbb{1}_p^l \times \mathbb{1}_p$ as the variable indicating whether pixel p is counted as part of the area eventually eligible for AOC l. We also define $N_{pt} = \sum_l \Delta_p^l \mathbb{1}_t^l$ as the number of distinct AOCs for which pixel p was eligible in year t.[18]

Finally, denoting by Σ_{it} the area in vineyards (under production or not) in department i in year t and by $P(i)$ the set of pixels in department i, we construct our main regressor as

$$s_{it}^k \equiv \frac{\sum_{p \in P(i)} \mathbb{1}_{N_{pt} \geq k}}{\max_t \Sigma_{it}}$$

which indicates the share of department i eligible for k or more AOCs as of year t. For six departments, our proxy of the acreage eligible for at least one AOC eventually becomes greater than the maximum area in vineyards[19]. In these cases, we set the share of eligible acreage to one.

[Table 1 about here.]

Table 1 shows summary statistics for a set of variables relevant to our analysis, including the shares s_{it}^1 and s_{it}^5 of vineyard acreage eligible for one or more and five or more AOCs, respectively. Figure 2 depicts the temporal rollout of AOC recognitions and, whenever available, the national area under AOC and the national

[18] In doing so, we do not double-count AOCs recognizing different colors of wine. For instance, if a parcel is eligible for producing both red and white AOC wine, we only count one AOC, the idea being that a given wine can only be sold under one color. As a result, the multiplicity of AOCs for a given parcel arises solely from hierarchical structure of the AOC system.

[19] These departments are Aisne, Aube, Charente, Charente-Maritime, Marne, and Haut-Rhin.

Figure 2 Temporal rollout of AOC recognitions and AOC production

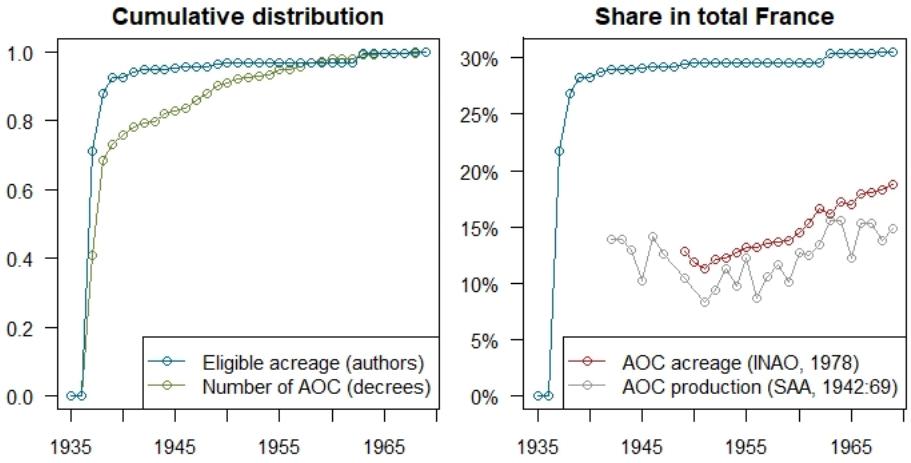

AOC wine production.[20]

3.2 Identification strategy

We exploit two sources of variation to identify the effects of the reform on the average wine price: variation in the exposure of a department to the reform (through its eligible share of vineyards) and variation in the timing of the decrees taken in application of the 1935 law. Most decrees were enacted during the years 1936 and 1937, although several were adopted later, including those pertaining to the Alsace region in 1962. Importantly, the reform affected wine-producing departments very unevenly: many had no AOC recognition, some full AOC recognition, and many had only a share of vineyards declared eligible for AOC status. This cross-sectional variation provides us with both an extensive and an intensive margin of treatment and allows us to control for common shocks to departmental wine prices through year fixed effects.

One legitimate concern when assessing the effect of a program or rule on outcomes is that implementation is not exogenous, i.e., rules happen to be implemented concurrently with other factors affecting the outcome. For instance, if appellation decrees happen to be enacted at the same time that demand factors, say expanding export markets, are affecting wine prices, then the effect of foreign demand might be mistakenly attributed to regulation if it happens to affect treated and untreated

[20]The area under AOC was obtained from estimates by INAO (INAO, 1978). AOC production is reported in *Statistique agricole annuelle*.

Figure 3 Definition of *vignobles*

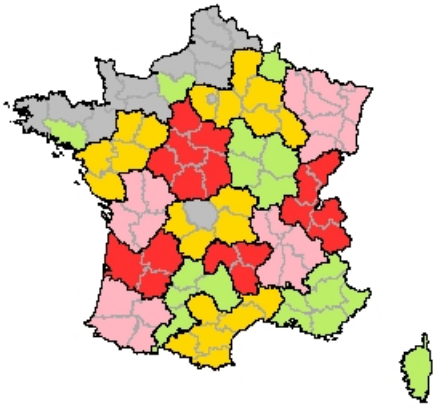

Note: Delineations in light gray represent departments. Delineations in black represent *vignobles*. Departments shaded in gray are excluded from the analysis because they produced little to no wine during the period.

departments differently. Our strategy to control for such potentially confounding factors is to further differentiate the year fixed effects by *vignoble*, that is, the broad geographical area that defines wines, such as "Loire" or "Midi." We define these *vignobles* so that each of them makes sense from a regional and viticultural standpoint. In fact, we largely follow the classification adopted by INAO, making sure that each *vignoble* is large enough to include at least a couple of departments, our cross-sectional units of analysis. Our dataset includes 16 *vignobles* and 76 departments, depicted in Figure 3.

Given the limited geographical span of our *vignobles*, we believe it is unlikely that remaining unobservables correlated with the AOC share within a *vignoble*-year could be confounding the effect of regulation. Controlling for *vignoble*-by-year fixed effects means that our identification relies on differences, within a *vignoble*, on the share of vineyards eligible for an AOC in a given year following the reform. Such differences arise from different shares of a department's territory being eligible for a given appellation and, to a lesser extent, from different dates of adoption of decrees for different appellations. For instance, if two departments in the same *vignoble* are only eligible for one and the same appellation, they will nonetheless participate in identification as long as they have different shares of vineyards eligible for that appellation. Conversely, if two departments in the same *vignoble* have the same share eligible, but this share relates to two distinct appellations with decrees taken at different dates, they will participate in identification as well. Assuming for a moment that decree adoption does cause an increase in wine prices, we would expect departments within a *vignoble* with larger shares of vineyards eligible to have higher price increments upon AOC recognition; we would also expect eligible departments within a *vignoble* to experience price increases sooner if their decrees are enacted sooner.

Formally, our preferred specification can be spelled out as follows:

$$\log p_{it} = \alpha_i + \gamma_{vt} + \boldsymbol{\beta}'\mathbf{s}_{it} + \boldsymbol{\delta}'\mathbf{x}_{it} + \epsilon_{it} \tag{4}$$

where i denotes a department, t denotes a year, v denotes the unique *vignoble* to which department i belongs, p_{it} is the average price of wine in department i in year t, α_i is a department fixed effect, γ_{vt} is a *vignoble*-by-year fixed effect, $\mathbf{x}_{it}$ is a vector of quantity controls, and $\mathbf{s}_{it}$ is a vector of treatment variables capturing the extent of AOC recognition in department i in year t. For instance, the vector $\mathbf{s}_{it}$

may include the share of a department's vineyard acreage eligible in year t for one or more and five or more controlled appellations. The vector $\boldsymbol{\beta}$ captures the effects of interest. Our identifying assumption is thus that within a *vignoble*, treated and untreated departments would have followed parallel price movements if not for the AOC reform. We provide support for this assumption before presenting our main results.

Note that irrespective of the structural and welfare interpretations of the coefficient on the eligible share discussed in Section 2, the estimate of $\boldsymbol{\beta}$ in Equation (4) has a clear reduced-form interpretation: it gives rate of increase of the departmental wine price with respect to the shares of vineyards declared eligible for k or more AOCs.

Our specification includes *vignoble*-by-year fixed effects. These fixed effects control non-parametrically for yearly shocks common to departments located in the same *vignoble*, notably those due to common weather shocks. On average, there are fewer than five departments in each *vignoble*. Our preferred standard errors assume that, conditional on these geographically differentiated yearly shocks and other included regressors, there is no residual correlation in errors across departments. Nonetheless, we allow for serial correlation across years within a department through the use of department-level clusters. We believe that this is important because we observe outcomes in many years, and treatment is correlated across years before and after the reform. We view department-clustered standard errors as conservative enough, particularly given the small number of departments within each *vignoble*. As a point of comparison, Jensen and Miller (2008) report standard errors clustered at the level of the treatment unit (a household) in a panel regression that includes county-by-year fixed effects, even though the number of households per county in their study is 100–150. Further, because we sample all departments in all *vignobles*, there is no sampling design justification for clustering at the *vignoble* level (there are no relevant *vignobles* absent from our data set that we wish to draw inference about). Instead, we view our sampling as occurring in the time dimension, in which case department-level clusters should be appropriate (Abadie et al., 2017).

For comparison purposes, we also report two other types of standard errors: (i) standard errors computed using the method of Conley (1999) adapted for panel data,[21] and (ii) standard errors clustered at the level of a *vignoble*. Note that unlike the

[21] The Conley errors are to spatial data what Newey-West errors are to time-series data. We apply

department- and *vignoble*-clustered standard errors, Conley errors do not account for serial correlation of the error term.

3.3 Results

Before we turn to our main regression results, we present simple suggestive evidence that AOC recognition positively affected the trajectory of wine prices at the department level.

3.3.1 Suggestive evidence

Figure 4 plots a time-series of average real wine prices across two categories of departments: those with high eventual AOC share (defined as those with an eligible share of AOC vineyards larger than 25% by 1969) and those with low eventual AOC share (defined as those with an eligible share lower than 2.5%). A few departments with intermediate share are not represented.

The figure suggests that before the reform the two categories of departments had very similar prices, while after the reform (whose implementation started with the first set of decrees published during 1936) average prices started to diverge between the two groups, with higher values in departments with high eventual AOC share. The figure admittedly provides visual evidence of the "parallel trends" assumption implicit in difference-in-differences designs.[22] What the figure does not capture, but our main regression will, is any differential price trends *within* the two broad categories defined here according to the AOC eligible share and the behavior of prices in departments with intermediate share (that is, the intensive margin of treatment along the AOC share dimension), as well as the fact that recognition did not happen simultaneously in all treated departments (the intensive margin of treatment along the time dimension).

Figure 5 depicts trends in real wine prices over the period 1910 to 1965 at the departmental level, using changes in 25-year averages from the endpoints of the period to compute the relative increase in price. It also depicts the share of vineyards eligible for AOC recognition by department as of 1965. Qualitatively, Figure 5 tells

the Newey-West weighting scheme to neighboring relationships when calculating our standard errors.

[22] Average prices in the nine departments with an intermediate eventual AOC share do not contradict this story: prices in those departments were consistently below those in non-AOC departments before the reform, and caught up after it.

Figure 4 Average real wine prices in AOC and non-AOC departments

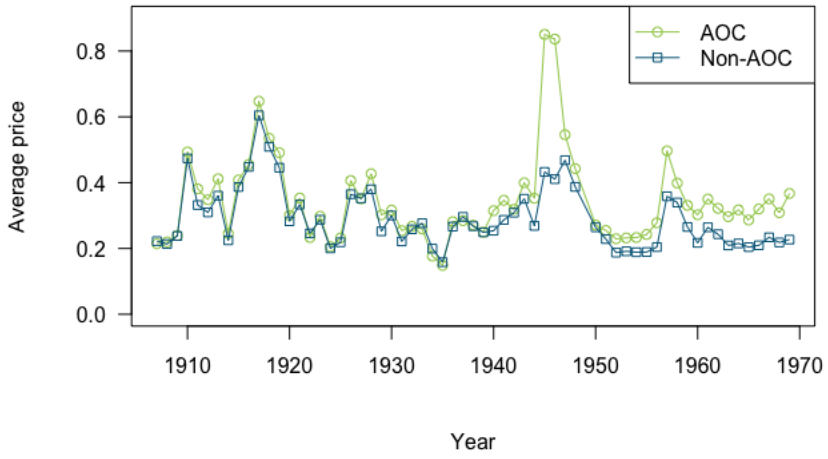

Note: Average real wine prices are calculated using production weights and conditioning on departments without missing data. Prices are deflated by a CPI. Production weights are constant over time and calculated as the average departmental wine production over the period. AOC departments (22) are departments with a 1969 share of vineyards eligible for AOC larger than 25%. Non-AOC departments (32) are departments with a 1969 share of vineyards eligible for AOC smaller than 2.5%. There are 9 departments with an intermediate share that are not represented.

Figure 5 Trends in departmental real wine prices over the period 1910-1965

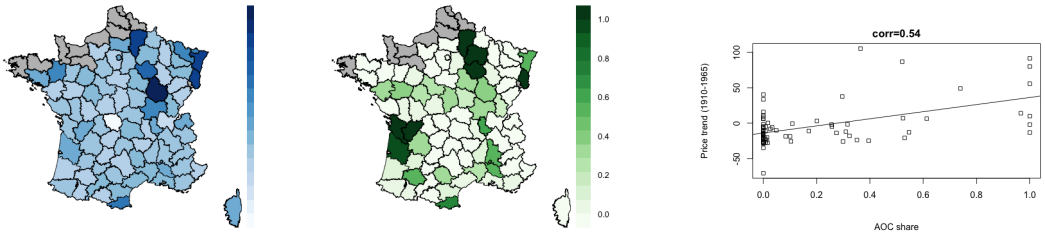

(a) Price trend (b) AOC share (c) Correlation

Note: Price trends are computed using changes in 25-year averages from the endpoints of the period and are expressed in relative terms. The share of vineyards eligible for AOC is calculated as of 1965. Gray departments: no data available.

a similar story as the previous figure: price trends over the period 1910-1965 appear to be stronger in departments with higher AOC shares.

[Table 2 about here.]

One may be worried that departments with eventually high shares of AOC recognition may have been on a steeper price trend for reasons unrelated to regulation. For instance, one could perhaps imagine that producers in departments with steeper price trends lobbied harder for AOC recognition. To investigate this possibility, we compare two simple price trend regressions based on different subsamples of years: 1907–1936 (pre-regulation) and 1927–1956 (pre-post-regulation), where price trends are computed using 10-year averages from the endpoints of each period and are expressed in relative terms. The results are reported in Table 2. Column (1) of the table reports the coefficient on the AOC eligible share (by 1956) from a regression of the price trend calculated over the period 1927–1956. Column (2) controls for *vignoble* to purge the regression of effects common to all departments located in the same wine region. In both columns, the coefficient on the AOC share is highly significant, suggesting that AOC eligibility had a positive effect on price trends, even after controlling for *vignoble* effects. In contrast, columns (3) and (4) show that if we focus on price trends during the pre-regulation period, the AOC share does not have any explanatory power, that is, eventual AOC eligibility (as of 1956) is irrelevant to explaining price trends prior to regulation. Finally, columns (5) and (6) show that AOC eligibility also had no clear effect on wine output, suggesting that the effects of regulation on price trends were not the result of a reduction in volumes.

3.3.2 Panel analysis

The results from the estimation of Equation (4) appear in Tables 3–5. Each table uses a different window of time to identify the effects of AOC recognition, from the widest (1907–1969, the entire data set) to the narrowest (1921–1950). Because it takes time for wines to (re-)establish a reputation, even after regulations have been adopted, we prefer to estimate effects using time windows that include a certain number of years after treatment.

[Table 3 about here.]

We do not necessarily expect coefficient estimates to be stable across periods. One basic reason is that as periods change, so does the set of appellations that are recognized in the sample. For instance, appellations in the Alsace region were recognized relatively late (1962). Because AOC recognition may cause different price increases in different regions, our coefficient estimate, which captures an average effect, may vary according to the period used. Despite this consideration, our results suggest a relatively consistent effect across time: AOC recognition did increase the price of wine, even after conditioning on quantity produced, by a non-negligible factor. In the largest sample, the regression with the richest set of fixed effects implies a 45% increase in the average wine price associated with full AOC recognition (column (4) of Table 3, using contemporaneous rather than lagged quantity as a control). The estimate is highly significant, no matter which standard error is used towards inference.

[Table 4 about here.]

Our tables report two different effects: that of the share of vineyards eligible for one AOC or more (regressor AOC Share1), and, in some regressions, that of the share of vineyards eligible for five AOCs or more (regressor AOC Share5). The share of vineyards eligible for one or more (resp. five or more) AOCs is 30.5% (resp. 2.8%) across all departments as of 1969, the last year in our sample. All tables suggest a sharp gradient with respect to the number of designations that a vineyard may claim. For instance, column (9) of Table 3 indicates that eligibility for one to four appellations increases the average wine price by 42%, while eligibility for five or more appellations increases it by 192% (the sum of the coefficients on the two share regressors). Again, these price effects are statistically significant, even when one considers narrower time windows.

[Table 5 about here.]

Our identifying assumption is that conditional on *vignoble*-by-year effects and quantity, there are no unobserved determinants of price correlated with the AOC eligible share. One could be concerned however that departments eligible for AOC recognition were on a different price trajectory than control departments. Although our *vignoble*-by-year fixed effects control for trends common to all departments within a *vignoble*, the relatively long period used makes it plausible that factors

that would have systematically propped up prices in treated departments after the reform, even if temporary, could be confounding the effect of regulation. For instance, an increasing taste of foreign markets for *bordeaux* wine happening after the reform could affect identification since the share eligible for the *bordeaux* AOC increased from zero to almost one within a single year.[23]

[Table 6 about here.]

In order to rule out such possibility, we first run a falsification test of the relationship between AOC recognition and average wine price using data from the pre-reform period. Specifically, we artificially set the share of AOC vineyards to its value fifteen years later. Since we begin counting AOC recognition the year after the decree is enacted, and the first decrees were enacted in 1936, from 1907 to 1921 (15 years) our AOC share remains equal to zero and the artificial treatment period goes from 1922 to 1936 (15 years). Results are shown in Table 6 and confirm that later AOC recognition was uncorrelated with price patterns before the actual treatment period began. The table also shows results obtained from models where the share of eligible vineyards is artificially set equal to its value 10 years later (columns (8) and (9)) or 20 years later (columns (7) and (8)). If anything, the results in columns (8) and (9) suggest a negative correlation between eventual AOC recognition and pre-reform price trends right before the reform, although statistical significance is dependent on which standard error is used. This relative erosion of wine prices in departments ultimately eligible for AOC recognition would be consistent with historical accounts of increasing abuses in the appellation wine market in the years leading to the reform (the "appellation scandal," see Capus (1947)).

[Table 7 about here.]

As a second piece of evidence that our estimated effects are not driven by unobserved correlated factors, we restrict the sample to the post-1937 period: by that date, the most important AOCs had already been defined, so that the residual variation in the AOC eligible share, conditional on the departmental fixed effects, comes from later rounds of AOC recognition, notably that of Alsatian wines. Results are displayed in Table 7. Although the point estimate is slightly smaller than in the full sample, the effect of the AOC share remains large (33% in column (4)) and

[23] In order to confound our effect, such a taste would have to be unrelated to AOC recognition, however.

statistically significant. For the estimated effects to be spurious, unobserved factors would thus need to be systematically correlated with waves of AOC recognition.

[Table 8 about here.]

Finally, we investigate whether the set of departments used as controls is a significant driver of our results. If unobserved factors unrelated to AOC recognition affected wine prices in AOC departments differently than in non-AOC departments within a *vignoble* after the reform, we would expect that removing non-AOC departments from the sample would change the estimated effect of AOC recognition. We thus remove all departments for which the eventual AOC share (by 1969) lies below 2.5%, as well as all departments with missing data. There are 45 such departments out of 76 used in the full sample, therefore this procedure removes more than half of the departments. Because departments with zero or very small eligible share are removed, identification now relies on comparisons of price changes across moderately and more intensively treated departments within the same *vignoble*, which are plausibly more similar to each other and less likely to be differentially affected by factors unrelated to AOC recognition after the reform. Results are displayed in Table 8 and show that the coefficient estimates on the AOC eligible shares are still statistically significant and of very similar magnitudes as those obtained using the entire sample of departments.[24] For instance, column (4) shows an average effect of AOC recognition equal to 43%, as opposed to 45% in the full sample.

Taken as a whole, these results clearly suggest that AOC recognition caused a sizeable appreciation in average wine prices at the department level. Our preferred estimate indicates a rate of increase of the average wine price with respect to the share eligible for AOC of 45%. This figure implies that in a department where 100% of vineyards became eligible for at least one AOC (like Gironde), the average wine price increased by 45%.

3.3.3 Ruling out alternative explanations

The results of Section 3.3.2 suggest a clear effect of AOC recognition on the departmental wine price. Whether the increase in wine price was indeed related to quality enhancements that failed to be incentivized prior to the reform remains to be established. Perhaps one of the biggest threats to identifying whether the

[24]The results still hold if we include departments with missing data in the regression.

AOC reform had any effect on the supply of quality is its potential for affecting the volumes of wine produced. There are at least two potential channels to consider: first, the reform could have reduced overall wine acreage and/or yields in regulated areas, and therefore the quantity of wine produced. Second, the reform could have reshuffled volumes of wines away from the appellation market into the ordinary wine market.

Acreage and yield effects

It is difficult to imagine how overall acreages could have been reduced by the reform because it did not force producers to uproot existing vineyards. One could easily expect, however, that maximum yields specified in many appellation decrees may have resulted in yield (and therefore production) reductions. In fact, we do not detect any negative effects of the reform on either acreage or yield in the data.

[Table 9 about here.]

Instead, regressions reported in Table 9 show that the share of AOC recognition had a positive (and in many cases statistically significant) effect on acreage planted, irrespective of the window of time selected for the regression. This is consistent with the view (confirmed by Figure 6) that wine acreage decreased more in non-AOC regions than in AOC regions over time. One potential explanation for the observed differential trends in acreage may be that the AOC reform increased the profitability of wine making in treated areas and therefore had an effect on the extensive margin of wine production. If this were the case, this additional supply of wine would be direct evidence of the increase in market size attributable to the resolution of the lemons problem, as described in Akerlof's original model of adverse selection whereby high-quality sellers do not sell in equilibrium. As appealing as this explanation may seem, we are unwilling to entirely attribute the relative acreage increase in AOC departments to the resolution of the lemons problem, the main reason being that AOC vineyards were largely exempt from restrictions on new plantings applicable to the ordinary wine market starting in 1953 (Humbert, 2011).

Table 9 also shows that there is no clear effect of AOC recognition on yield. Estimates are small with fluctuating signs, and generally not statistically significant. Although many AOC decrees specify maximum yields, it thus appears that reducing yield was not the principal channel through which quality improvements were achieved. In a way, the absence of a negative effect of AOC recognition on average

Figure 6 Area in vineyards in AOC and non-AOC departments

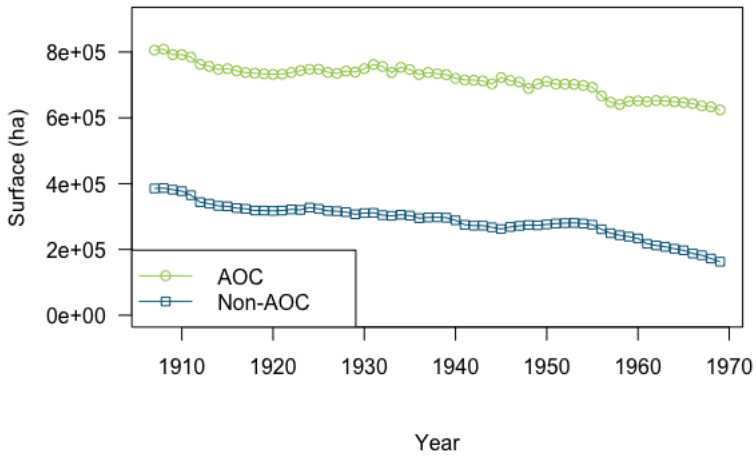

Note: Areas excludes departments with missing data. AOC departments (22) are departments with a 1969 share of vineyards eligible for AOC larger than 25%. Non-AOC departments (35) are departments with a 1969 share of vineyards eligible for AOC smaller than 2.5%.

departmental yield is consistent with available information. For instance, the large regional *bordeaux* and *bourgogne* AOCs, when created, imposed a maximum yield of 50 and 45 hl/ha, respectively. The average yield over the ten years prior to regulation were 32 hl/ha in Gironde—the department where *bordeaux* is produced—and 28 hl/ha in the departments covered by the *bourgogne* AOC.

Reshuffling effect

The second main effect that the reform could potentially have had on wine quantities is redistributive. It is conceivable that some wines that used to be sold under appellation before the reform were later denied the use of the AOC and thus had to be sold either under a less prestigious name, if available, or as ordinary wine (an effect known as *déclassement*). Initially, the "double appellation" regime allowed wines to be sold under a plain appellation *of the same name* as a recognized AOC, as long as they met the geographical requirements associated with the use of the denomination (that is, the rules already applicable prior to the reform). During that initial period, it is thus unlikely that the reform led to significant *déclassement* as wines could continue to be marketed pursuant to the old rules. This tolerance was formally abolished in 1942 however, which implies that wines not meeting the more

comprehensive AOC requirements could no longer claim the appellation, even if they originated in an eligible region. It is therefore likely that the reform ultimately resulted in the reallocation (or "reshuffling") of some volume of wine away from the appellation market into the ordinary wine market.

If consumers are homogenous with respect to their taste for quality, such movements from one segment of the wine market to another should leave the average price of wine at the department level unchanged as consumers update their valuations of ordinary and appellation wines based on the average quality present in each market segment. In that case, the reshuffling effect should not confound our finding that average prices increased due to AOC recognition through an increase in the quality of AOC wines.

However, as we formally show in Appendix A.2, if consumers are heterogenous with respect to their taste for quality, average price could increase as a result of the *déclassement* of wines alone, *even if the reform has no effect on quality*. In that case, we show that welfare increases as well through a reallocation effect, as higher-quality wines are more selectively chosen by quality-valuing consumers. However, the relationship between the relative increase in price and the relative increase in welfare is much less straightforward than in the case where the price increase is due to an increase in the quality of AOC wines. In addition, while reshuffling *could* cause a price increase, it could also cause a price decrease, so the net contribution of the reshuffling effect to our overall effect remains ambiguous.

In order to assess whether the positive and significant effect of AOC recognition on the average wine price found above is driven by reshuffling, rather than quality improvements, we leverage additional data to assess the extent to which appellation wines were forced into the ordinary wine market after the reform.

France's *Journal Officiel* for the year 1936 reports the volumes of wines declared under appellation in the year 1935, by department. We pair these data with production data from the *Statistique agricole annuelle* to construct a measure of the volume share of wines sold under (plain) appellation right before the reform. (We only have data for 1935, but this year, together with 1934, had historically high shares of appellation wines because producers were trying to escape the constraints of the new *Statut Viticole* applicable to ordinary wines (Capus, 1947). Therefore this approach overestimates the share of wines sold under appellation over the entire pre-reform period.) We compare these departmental appellation shares to the average departmental share of wines sold under either AOC or plain appellation in the

post-1942 years, after the "double appellation" regime ended.[25] We then identify departments for which the share decreased by more than 5% (that is, for instance, a share going from 40% to 37% or less), and exclude these from the sample if the share of vineyards eligible for AOC during the post-reform years (our regressor of interest) was nonzero. The idea is that in departments with AOC recognition, a post-reform appellation share lower than the pre-reform appellation share could have plausibly been caused by *déclassement*. The 5% tolerance is meant to account for variations in volumes due to weather, as well as for the fact that the 1935 share likely overstates the share that prevailed during the broader pre-reform period. As such, we believe that this approach is conservative.

[Table 10 about here.]

Table 10 shows estimation results for two different sample periods: 1921–1950 and 1907–1969. In each case, departments included in the regression are selected based on a post-reform appellation share computed over years present in the sample, that is, 1942–1950 or 1942–1969. For comparison purposes, the table also reports coefficient estimates for the full sample of departments over the same time periods. The decrease in sample size once we remove departments with a decline in the appellation share indicates that our procedure removes a sizable portion of the initial sample. Indeed, key departments are excluded, including, for the 1921–1950 sample, the *bordeaux* region (Gironde department), most of the *bourgogne* and *côtes-du-rhône* regions, as well as a large share of the *champagne* region. (Fewer departments are excluded when looking at the longer time period as the share of wines sold under AOC rose over the post-reform period.) Nonetheless, irrespective of the period considered, the estimates in Table 10 suggest that if anything, the estimated effect is stronger when removing those departments with plausible reshuffling. Therefore, it is unlikely that our estimated positive effect is driven by the pure reallocation of wines across market segments following the reform.

Linearity of the AOC eligibility effect

Our regression model in Equation (4) assumes a linear total effect of AOC eligibility on the log-price of wine. To test whether this linearity assumption is justified by the data, we estimate flexible models that nest the linear model and perform post-estimation specification tests of the linearity assumption. Specifically,

[25] This information is available starting in 1942 from France's *Statistique agricole annuelle*.

we estimate a model where the eligible share enters quadratically on the right-hand side, and a model whereby we discretize the AOC share using three categories: $0 < s_{it} \leq 0.05$, $0.05 < s_{it} \leq 0.40$, and $0.40 < s_{it} \leq 1$. The cut-off points were chosen in order to ensure a sufficient number of observations falling in each category. The categorical model estimates coefficients on dummy variables for each category. For the quadratic model, the linearity test is simply the t-test on the quadratic coefficient. For the model with categorical share variables, the linearity test is a set of linear restrictions involving the mean shares within each category. We report tests statistics constructed using the department-clustered variance-covariance matrix.

[Table 11 about here.]

Table 11 shows that the linearity assumption cannot be rejected at standard levels of statistical significance for most specifications, the exception being the test based on a quadratic model with coarser year fixed effects (as opposed to *vignoble*-by-year fixed effects), and only when using the less conservative spatial-robust standard errors that ignore serial correlation.

Other robustness checks

Tables 12 and 13 provide results for samples that exclude selected years or selected departments. Table 12 investigates the robustness of our estimated effects to the removal of (i) the German occupation years 1940–1945, (ii) the post-war years 1945–1947, during which there was a sharp increase in wine prices in AOC departments (e.g., Gironde), and (iii) the years 1941–1947, during which ordinary wines as well as certain AOC wines were subject to administered prices (a regime known as *taxation*) and requisition.[26] Table 13 shows results for samples that exclude (i) the four departments of the Champagne region, which had production requirements enacted (without official control) as soon as 1927, and (ii) Gironde, the department where *bordeaux* wines are exclusively produced (and the home department of Joseph Capus, the assemblyman who promoted the 1935 law).

[Table 12 about here.]

Results obtained when removing selected years are in line with those for the full sample, irrespective of the window of time omitted from the sample.

[26] Taxation lasted after the liberation and ended with a decree of 4 September 1947, see Milhau (1949) and Humbert (2011).

[Table 13 about here.]

Results without Champagne departments make sense. Unlike other wine regions, Champagne does not have sub-regional appellations, therefore the vast majority of eligible vineyards in Champagne departments are only eligible for one appellation, *champagne*. Despite this fact, *champagne* is perhaps the most prestigious of all wine appellations and the one that commands the highest prices per hectoliter. To the extent that *champagne* benefited relatively more than other appellations from AOC recognition, which is plausible, its effect would solely be captured by the AOC Share1 regressor. Including Champagne departments in estimation would then tend to pull the estimate on the AOC Share 1 towards a slightly higher value than when these departments are omitted. As the effect on AOC Share1 is being pulled upwards by Champagne departments, the coefficient on AOC Share5 is decreased as a larger share of the effect is already been captured by AOC Share1.

Finally, results without *bordeaux* wines are very similar to the results for the full sample, suggesting that the reform was effective at promoting quality well outside of Gironde.

3.3.4 Heterogeneity analysis

Heterogeneity across time

We would expect that the new regulatory apparatus introduced by the 1935 reform took some time to generate discernable effects on the French wine market. Such expectation is grounded on both supply-side and demand-side considerations. First, some of the production requirements introduced in AOC decrees could have required sizable upfront investments, notably grape varietal requirements for vineyards that were previously planted in unapproved varieties. This could imply a lag between AOC recognition and the actual increase in wine quality. Second, even if producers were able to adjust rapidly to the newly mandated requirements, it could have taken time for consumers to update their beliefs regarding the underlying quality of AOC wines as wine quality is subject to the vagaries of climate, making quality assessment difficult within just a few years. Finally, some eligible but perhaps high-cost producers may have chosen to delay changes in production practices until they could better assess the extent to which the market would reward their investments towards quality. This last point would be consistent with the AOC eligibility and acreage trends shown in Figure 2, which clearly suggest

that the compliance rate for eligible producers (i.e., the AOC acreage relative to the AOC eligible acreage) increased over time.

[Table 14 about here.]

In order to test the proposition that the effect of AOC eligibility on wine prices increased over time, we estimate a variant of Equation (4) where we interact the AOC eligible share with a linear time trend. Results are shown in Table 14 and suggest that the effect of AOC recognition of the wine price started at about zero in 1937, the year immediately following the first round of AOC decrees, and then increased at a rate of about 2.6% per year in our preferred specification (column (4)). The slope coefficient is precisely estimated.[27]

Heterogeneity across space

Given the diversity of wines across France's *vignobles*, one may expect that the effect of the reform on the wine price could differ geographically. Indeed, reputable regions differ in their propensity to produce red vs. white wines, still vs. sparkling wines, but also in the way that appellations themselves are organized, with many idiosyncratic factors like the use of *châteaux* to identify prestigious wines in Gironde, the widespread use of municipal appellations in Burgundy, etc.

In order to investigate heterogeneity in the effect of AOC recognition across *vignobles*, we estimate a variant of Equation (4) where we interact the AOC eligible share with regional dummy variables. In order to keep the model tractable and well identified under our rich set of geographically differentiated year effects, we choose to group *vignobles* according to geographical proximity. This leads us to defining five broad groups of *vignobles*, represented in panel (a) of Figure 7. The spatial heterogeneity of the effect is represented in panel (b) of Figure 7. Although all coefficients are positive, two out of five regions (regions 1 and 4) show smaller effects with statistical significance contingent on the choice of standard error, whereas the other three regions show large and significant effects, in excess of 50%.

3.3.5 Effect of the reform on gross welfare and implied price premium

At the end of our study period, the share of vineyards eligible for at least one AOC was 30.5% across all departments. Together with an estimated effect of AOC recognition of 45%, this figure implies a relative increase in gross welfare in the French

[27] We also estimated a model with a quadratic, rather than linear, trend but the coefficient on the quadratic interaction term was not statistically significant.

Figure 7 Heterogeneity of the effect across space

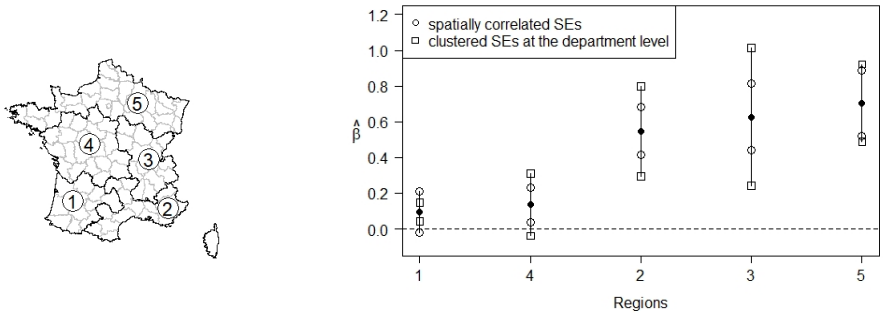

(a) *Vignoble* groupings (b) Estimated effects

Note: These estimates are obtained using the specification with Year×*vignoble* fixed effects and demand flexibilities with respect to contemporaneous production differentiated by *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years. Price data is missing for the year 1949. The bounds of the intervals are +/- $\Phi^{-1}(0.975)$ times the chosen standard error of the estimate, with Φ being the cumulative distribution function of the standard normal.

wine market (inclusive of ordinary wine) of about 14%, ignoring potential increases of production costs for wines sold under AOC after the reform. Furthermore, the increasing effect exhibited in Table 14 implies even larger gains at the end of the period.

As indirect evidence that the additional costs of supplying AOC wines were generally smaller than the price premium, we collected supplementary data to evaluate the share of eligible producers who found it profitable to produce AOC wines. A technical report by INAO gives the shares of total vineyards cultivated for the production of AOC wines from 1949 through 1969 (INAO, 1978). Over that period, the share of total acreage producing AOC wines increased from 12.8% to 18.8%, while we estimate that the share of the eligible area relative to the area in vineyards increased from 30.8% to 34.6%.[28] These figures imply that the share of eligible area where AOC wine production was found to be profitable grew from 41.6% to 54.3%. This widespread and increasing adoption of AOC denominations suggests that the price premium for AOC wines exceeded the additional production

[28]The figure of 34.6% in 1969 differs from the 30.5% mentioned above because the share is computed relative to the area in vineyards *in the current year*, as opposed to the *maximum* of the area in vineyards across all sample years.

cost for the inframarginal units supplied, and thus that the reform was welfare-enhancing.

Retrieving the price premium from our estimate is less straightforward. Because not all eligible parcels claim an AOC, our estimated coefficient $\hat{\beta}$ on the share of vineyards eligible for AOC recognition underestimates the effect of the reform on the price of wines actually sold under an AOC. However, using additional assumptions, our data allow us to size this effect. Using the notation of Section 2, the average price can be written as $p_m = p_0 + s_1(p_1 - p_0) + s_2(p_2 - p_1)$ with s_2 denoting the share of production sold under AOC. Since our empirical model regresses $\log p_m$ on σ_2, the share of acreage eligible for AOC, it is useful to express p_m as a function of σ_2 rather than s_2. Let κ be the share of eligible acreage actually cultivated for AOC wine production, and let y_1 and y_m denote the appellation yield and the average yield, respectively.[29] We then have $s_2 = \sigma_2 \kappa \frac{y_1}{y_m}$. Therefore,

$$p_m = p_0 + s_1(p_1 - p_0) + \sigma_2 \kappa \frac{y_1}{y_m}(p_2 - p_1).$$

Our empirical model uses the share of acreage eligible and therefore estimates $\hat{\beta} = \frac{\partial \log p_m}{\partial \sigma_2} = \frac{p_2 - p_1}{p_m} \kappa \frac{y_1}{y_m} = \frac{p_2 - p_1}{p_m} \frac{s_2}{\sigma_2}$. Therefore,

$$\frac{p_2 - p_1}{p_m} = \hat{\beta} \frac{\sigma_2}{s_2}.$$

Between the years 1942 and 1969 (excluding 1948 and 1950) the agricultural yearbook reports the volume of wines sold under an AOC for each department, allowing us to compute the share s_2 for each department and each year over this period. Using our proxies of σ_2, we calculate a geometric mean ratio for $\frac{\sigma_2}{s_2}$ equal to 2.34 over all treated departments and all years where s_2 is observed. Assuming that this mean holds over the whole study period, the price effect relative to the mean price, $\frac{p_2 - p_1}{p_m}$, is estimated to be 105%. This figure implies an average appreciation for AOC wines roughly equal to the size of the average wine price.

[29] As in Section 2, it is assumed that yields are unaffected by the reform. In addition, here we assume that κ and $p_2 - p_1$ do not depend on σ_2.

4 Discussion

This article provides empirical evidence that the quality of wines sold under appellation prior to a pioneering 1935 law was below the social optimum, and that the reform allowed producers to profitably adopt quality-enhancing practices. Using average wine price data at the department level, we show that trends between the pre- and post-reform period are strongly correlated with the eventual AOC share. Using a panel approach with fixed effects, we estimate that the price of wines ultimately sold under an AOC increased by a value roughly equal to the average wine price.

In order to interpret this remarkable price increase as stemming from an increase in wine quality, we first rule out that AOC recognition negatively affected the quantity of wine produced in treated departments. We then show that the reshuffling of previous appellation wines towards the ordinary wine market did not contribute to the observed price increase. In the end, the best explanation for the large and significant price increase in treated departments is that the reform had the intended effect: it provided incentives to atomistic producers to supply costly quality valued by consumers.

Although treated departments are different from control departments in the sense that they benefit from natural factors that are conducive to producing higher-quality wine, we provide several arguments in support of the parallel trends assumption necessary for a causal interpretation of our results. First, we show that eventual AOC status was completely uncorrelated with price patterns during the pre-reform period. Second, we show that the estimated effect survives the elimination of the pre-reform years from estimation. That is, exploiting only the period after the first wave of AOC recognitions yields a comparable estimate of the effect of AOC recognition. Finally, the estimated effect is robust to eliminating from the sample the set of control departments, defined as those with incomplete data or an eventual AOC share lower than 2.5%. That is, exploiting only the intensive margin of AOC eligibility to identify the effect does not change our overall estimate.

We provide a simple theoretical framework to underscore the welfare implications of our estimated price effect. If consumers are homogenous, then the relative price increase can be directly interpreted as a relative increase in consumers' average valuation of wine, or gross welfare. If consumers are heterogenous, then the estimated effect still constitutes a lower bound to the valuation increase relative to

market value.

More work is needed to assess the effect of AOC adoption on production costs, but the evidence presented here, together with the observation that AOC was widely adopted wherever available, suggests that the reform was welfare-enhancing. Ignoring increases in production costs, our estimate implies that welfare increased by 14% in the French wine market due to the reform.

References

Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. (2017). When should you adjust standard errors for clustering? Technical report, National Bureau of Economic Research. NBER Working Paper No. 24003.

Akerlof, G. A. (1970). The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.

Ali, H. H., Lecocq, S., and Visser, M. (2008). The Impact of Gurus: Parker Grades and *en Primeur* Wine Prices. *The Economic Journal*, 118(529):F158–F173.

Andrabi, T., Das, J., and Khwaja, A. I. (2017). Report cards: The impact of providing school and child test scores on educational markets. *American Economic Review*, 107(6):1535–1563.

Ashenfelter, O. (2008). Predicting the quality and prices of Bordeaux wine. *The Economic Journal*, 118(529):F174–F184.

Bond, E. W. (1982). A Direct Test of the "Lemons" Model: The Market for Used Pickup Trucks. *The American Economic Review*, 72(4):836–840.

Bundorf, M. K., Levin, J., and Mahoney, N. (2012). Pricing and Welfare in Health Plan Choice. *American Economic Review*, 102(7):3214–48.

Capus, J. (1947). *La Genèse des Appellations Contrôlées*. INAO, France.

Cardon, J. H. and Hendel, I. (2001). Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey. *The RAND Journal of Economics*, 32(3):408–427.

Cawley, J. and Philipson, T. (1999). An Empirical Examination of Information Barriers to Trade in Insurance. *American Economic Review*, 89(4):827–846.

Chiappori, P.-A. and Salanié, B. (2000). Testing for Asymmetric Information in Insurance Markets. *Journal of Political Economy*, 108(1):56–78.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1):1–45.

Council of the European Communities (1970). Regulation (EEC) No 816/70 of the Council of 28 April 1970 laying down additional provisions for the common organisation of the market in wine. *Official Journal of the European Communities*, No L 99/1 (5 May 1970):234–251.

Crozet, M., Head, K., and Mayer, T. (2012). Quality sorting and trade: Firm-level evidence for French wine. *The Review of Economic Studies*, 79(2):609–644.

Cutler, D. M. and Reber, S. J. (1998). Paying for health insurance: the trade-off between competition and adverse selection. *The Quarterly Journal of Economics*, 113(2):433–466.

Dranove, D., Kessler, D., McClellan, M., and Satterthwaite, M. (2003). Is more information better? The effects of ?report cards? on health care providers. *Journal of political Economy*, 111(3):555–588.

Duflo, E., Greenstone, M., Pande, R., and Ryan, N. (2013). Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from india. *The Quarterly Journal of Economics*, 128(4):1499–1545.

Einav, L. and Finkelstein, A. (2011). Selection in Insurance Markets: Theory and Empirics in Pictures. *Journal of Economic Perspectives*, 25(1):115–38.

Einav, L., Finkelstein, A., and Cullen, M. R. (2010). Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics*, 125(3):877–921.

Finkelstein, A. and Poterba, J. (2004). Evidence on Adverse Selection: Equilibrium Signaling and Cross-Subsidization in the Insurance Market. *Journal of Political Economy*, 112(1):183–208.

Genesove, D. (1993). Adverse Selection in the Wholesale Used Car Market. *Journal of Political Economy*, 101(4):644–65.

Hackmann, M. B., Kolstad, J. T., and Kowalski, A. E. (2015). Adverse Selection and an Individual Mandate: When Theory Meets Practice. *American Economic Review*, 105(3):1030–66.

Haeck, C., Meloni, G., and Swinnen, J. (2018). The Value of Terroir. A historical analysis of Bordeaux and Champagne geographical indications. LICOS Discussion Papers 40818, LICOS - Centre for Institutions and Economic Performance, KU Leuven.

Handel, B. R. (2013). Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. *American Economic Review*, 103(7):2643–2682.

Humbert, F. (2011). *L'INAO, de ses origines à la fin des années 1960: genèse et évolutions du système des vins d'AOC*. PhD thesis, Université de Bourgogne, France.

INAO (1978). Bulletin de l'INAO. Nouvelle série, numéro 9, 3ème trimestre.

Jensen, R. (2007). The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector. *The Quarterly Journal of Economics*, 122(3):879–924.

Jensen, R. and Miller, N. H. (2018). Market Integration, Demand, and the Growth of Firms: Evidence from a Natural Experiment in India. *American Economic Review*, 108(12):3583–3625.

Jensen, R. T. and Miller, N. H. (2008). Giffen behavior and subsistence consumption. *American Economic Review*, 98(4):1553–77.

Jin, G. Z. and Leslie, P. (2003). The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2):409–451.

Kim, J.-C. (1985). The market for "lemons" reconsidered: A model of the used car market with asymmetric information. *The American Economic Review*, 75(4):836–843.

Lewis, G. (2011). Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors. *American Economic Review*, 101(4):1535–1546.

Milhau, J. (1949). L'évolution du prix du vin au cours de la campagne 1947–1948. *Études et conjoncture – Union française / Économie française*, 2:57–62.

Ministère de l'agriculture (1937). Circulaire aux inspecteurs et agents de la répression des fraudes relative aux appellations d'origine contrôlées. *Journal Officiel de la République Française*, 28 Février 1937:2569–2571.

Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18(2):301–317.

Panhans, M. (2019). Adverse Selection in ACA Exchange Markets: Evidence from Colorado. *American Economic Journal: Applied Economics*, 11(2):1–36.

Puelz, R. and Snow, A. (1994). Evidence on Adverse Selection: Equilibrium Signaling and Cross-Subsidization in the Insurance Market. *Journal of Political Economy*, 102(2):236–257.

For Online Publication: Appendices

A Models with heterogenous consumers

We consider a Mussa-Rosen model of vertical differentiation (Mussa and Rosen, 1978) whereby consumer tastes are parameterized by an index $\theta \in [0,1]$ and $F(\theta)$ denotes the c.d.f. of θ. Each consumer consumes at most one unit of wine. The mass of consumers is set to $M > Q$, where Q denotes the fixed quantity of wine produced, including ordinary and appellation wine. Therefore, some consumers are not served in equilibrium. When consuming wine of quality μ sold at price p, a consumer of type θ enjoys utility $U^\theta(\mu, p) = \bar{u} + \theta\mu - p$, where $\bar{u} > 0$, and zero if the consumer purchases nothing.

A.1 Model with quality enhancement

This model is an alternative to the model presented in Section 2 whereby consumers are allowed to differ in their taste for quality. As in the main text, we assume that the effect of the reform is to increase the quality of a share of wine production previously sold under an appellation.

Wine quality is denoted $\mu_0 = 0$ for ordinary wine, $\mu_1 \geq \mu_0$ for an appellation that does not end up being recognized as AOC (i.e., a plain appellation), and $\mu_2 \geq \mu_1$ for an AOC.

Given that $M > Q$, the equilibrium price of ordinary wine must be equal to $\bar{u}$ ($p_0 = \bar{u}$) so that low-θ consumers are indifferent between purchasing nothing and purchasing ordinary wine.

Denote by $\tilde{\theta}$ the index of the consumer indifferent between purchasing ordinary wine and appellation wine. It must be that $\tilde{\theta} = \frac{p_1 - p_0}{\mu_1} = \frac{p_1 - \bar{u}}{\mu_1}$. Similarly, denoting by $\hat{\theta}$ the index of the consumer indifferent between purchasing plain appellation wine and AOC wine, we have $\hat{\theta} = \frac{p_2 - p_1}{\mu_2 - \mu_1}$. Market clearing implies that $M \int_{\tilde{\theta}}^{1} dF(\theta) = Qs_1$ and $M \int_{\hat{\theta}}^{1} dF(\theta) = Qs_2$ under full information. Under asymmetric information, all appellation wine has quality μ_1 and only the first market-clearing condition applies.

The relationship $M \int_{\tilde{\theta}}^{1} dF(\theta) = Qs_1$ determines $\tilde{\theta}$ given the exogenous values of Q, M, and s_1, and given $\tilde{\theta} = \frac{p_1 - \bar{u}}{\mu_1}$ it further determines p_1, which is then independent of the information regime. Similarly, the relationships $M \int_{\hat{\theta}}^{1} dF(\theta) = Qs_2$ and $\hat{\theta} = \frac{p_2 - p_1}{\mu_2 - \mu_1}$ determine $\hat{\theta}$ and p_2 under full information.

The increase in gross welfare (ignoring the additional costs of quality provision) when moving from the asymmetric to the full information scenario is simply the added gross utility of consumers with value index between $\hat{\theta}$ and 1, that is, those

Figure 8 Consumer utility and welfare under full information and asymmetric information

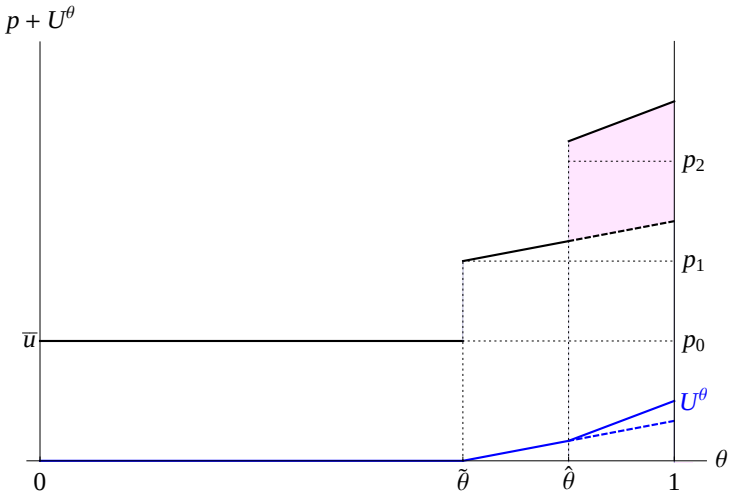

Note: Full information equilibria are represented with solid lines. Dashed lines represent outcomes, under asymmetric information, for consumers purchasing controlled appellation wine under full information.

with the highest tastes for quality who end up purchasing AOC wine:

$$\begin{aligned}
\Delta GW &= M \int_{\hat{\theta}}^{1} (\mu_2 - \mu_1)\theta dF(\theta) \\
&= M(\mu_2 - \mu_1) \int_{\hat{\theta}}^{1} \theta dF(\theta) \\
&= Qs_2(\mu_2 - \mu_1) \times \frac{\int_{\hat{\theta}}^{1} \theta dF(\theta)}{\int_{\hat{\theta}}^{1} dF(\theta)} \\
&= Qs_2(p_2 - p_1) \times \frac{\int_{\hat{\theta}}^{1} \theta dF(\theta)}{\hat{\theta} \int_{\hat{\theta}}^{1} dF(\theta)}.
\end{aligned}$$

Since $\frac{\int_{\hat{\theta}}^{1} \theta dF(\theta)}{\hat{\theta} \int_{\hat{\theta}}^{1} dF(\theta)} > 1$, it is clear that the gross welfare measure $Qs_2(p_2 - p_1)$ that holds with perfectly elastic demands (see Section 2) represents a lower bound to the gross welfare increase in the more general model.

Figure 8 illustrates the gross welfare calculation in the special case where $s_1 = 0.50$, $s_2 = 0.25$, $M = \frac{3}{2}Q$ and consumer taste parameters are uniformly distributed on $[0, 1]$. In this case, market clearing implies that $\tilde{\theta} = \frac{2}{3}$ and $\hat{\theta} = \frac{5}{6}$. Setting $\bar{u} = 1$, $\mu_1 = 1$, and $\mu_2 = 2$, we obtain the equilibrium prices $p_1 = \frac{5}{3}$ and $p_2 = \frac{5}{2}$. In the

figure, blue lines are used to represent consumer utility (net of the price paid) as a function of the taste parameter. Black lines represent the resulting gross welfare (ignoring supply costs) in equilibrium. Dashed lines depict utility and gross welfare for high θ consumers under asymmetric information. The shaded area represents the increase in gross welfare resulting from regulation.

A.2 Model with exogenous quality

A competing explanation as to why the average price rises after the reform (besides an increase in quality) is that the quantity of wine sold under appellation decreases as some wines are subject to *déclassement* (keeping constant the total quantity of wine sold). Indeed, wines previously sold under an appellation that did not meet the requirements for the AOC had to be sold either under a less prestigious appellation, if available, or as ordinary wine. If massive quantities of wines previously sold under appellation were redirected to the ordinary wine market due to the reform, the average price could change without any change in quality.

Here we thus assume that wine quality (and quantity) are fixed. We denote by $\mu_0 = 0$ the quality of ordinary wines, and by μ_1 the intrinsic quality of "true" appellation wines. The share of true appellation wines is s_1, but some of the ordinary wine is sold under appellation. The share of wine sold under appellation is thus $s_2 > s_1$. Therefore, the average quality of appellation wine is $\bar{\mu}_1 = \frac{\mu_1 s_1 + \mu_0 (s_2 - s_1)}{s_2} = \frac{\mu_1 s_1}{s_2}$. We assume the reform reduces the share of appellation wines by removing some of the low-quality wine from the appellation and forcing it to be sold as ordinary wine (its true quality).

At a market equilibrium, it must be that $p_0 = \bar{u}$ so that low-θ consumers are indifferent between consuming ordinary wine and consuming nothing. In addition, the index of the consumer who is indifferent between ordinary and appellation wine must satisfy $\bar{u} - p_0 = \bar{u} + \tilde{\theta}\bar{\mu}_1 - p_1$, which implies that $p_1 = \bar{u} + \tilde{\theta}\frac{\mu_1 s_1}{s_2}$. Market-clearing further implies that $M \int_{\tilde{\theta}}^{1} dF(\theta) = Q s_2$, which implicitly defines $\tilde{\theta}$ as a function of s_2. The average price of wine is then

$$\begin{aligned} p_m &= p_0(1 - s_2) + p_1 s_2 \\ &= \bar{u} + \mu_1 s_1 \tilde{\theta}(s_2). \end{aligned}$$

It is clear that $\tilde{\theta}$ decreases with s_2, so if the reform decreases s_2 to $s_2' < s_2$, we would expect $\tilde{\theta}$ to increase and the average price to increase. Note that this result critically depends on the presence of consumer heterogeneity: if all consumers are the same and wine quality does not change, then average price (and welfare) do not change in equilibrium, even if there is a redistribution of volumes towards the ordinary wine category. Let us now show that welfare also increases (in this case there is no reason to distinguish gross from net welfare as we assume away any quality enhancement). Denoting by $\tilde{\theta}'$ the index of the indifferent consumer after

the reform, we have

$$\begin{aligned}\Delta W &= -M\int_{\tilde{\theta}}^{\tilde{\theta}'}\theta\frac{\mu_1 s_1}{s_2}dF(\theta) + M\int_{\tilde{\theta}'}^{1}\theta\mu_1 s_1\left(\frac{1}{s_2'}-\frac{1}{s_2}\right)dF(\theta)\\ &= Q\mu_1 s_1\left[\frac{\int_{\tilde{\theta}'}^{1}\theta dF(\theta)}{\int_{\tilde{\theta}'}^{1}dF(\theta)} - \frac{\int_{\tilde{\theta}}^{1}\theta dF(\theta)}{\int_{\tilde{\theta}}^{1}dF(\theta)}\right]\\ &> 0\end{aligned}$$

while the change in price is simply $\Delta p_m = \mu_1 s_1\left(\tilde{\theta}' - \tilde{\theta}\right) > 0$. Therefore, in this case both price and welfare increase. But without further restrictions on the cumulative density function $F(\theta)$, it is not possible to determine whether the observed relative price increase attributable to the reform under- or -overstates the associated change in welfare, although both have the same sign. Also note that the only source of the welfare increase here is allocative efficiency as lower-quality (ordinary) wine is being redirected towards low-θ consumers.

Of course, the reason behind the welfare increase here is that the reform is able to sort out low-quality wine from the appellation market and redirect it to the ordinary wine market, so information improves. What if instead the reform arbitrarily redirects high-quality wine towards to ordinary wine market? To investigate this scenario, we now assume that before the reform $s_2 = s_1$ but after the reform $s_2' < s_1$, that is, only a portion of the high-quality wine has a right to the appellation. Average quality in the ordinary wine market is $\mu_0 = 0$ before the reform and $\bar{\mu}_0 = \frac{\mu_0(1-s_1)+\mu_1(s_1-s_2')}{1-s_2'} = \frac{\mu_1(s_1-s_2')}{1-s_2'}$ after the reform.

Market-clearing in the wine market determines the taste parameter of the lowest-θ wine consumer, $\underline{\theta}$, through the equality $M\int_{\underline{\theta}}^{1}dF(\theta) = Q$. Market-clearing in the appellation market determines the taste parameter of the consumer indifferent between ordinary and appellation wine: $M\int_{\tilde{\theta}}^{1}dF(\theta) = Qs_1$, and similarly after the reform: $M\int_{\tilde{\theta}'}^{1}dF(\theta) = Qs_2'$, so that $\tilde{\theta}' > \tilde{\theta}$. Consumers with taste parameter above $\tilde{\theta}'$ are consuming high-quality wine before and after the reform. Consumers with taste parameter between $\tilde{\theta}$ and $\tilde{\theta}'$ switch from high-quality wine to a mixture or low- and high-quality wine of quality $\bar{\mu}_0$. Low-θ wine consumers switch from

low-quality wine to that same wine mixture. Therefore, the effect on net welfare is

$$\begin{aligned}
\Delta W &= M\int_{\underline{\theta}}^{\tilde{\theta}} \theta \frac{\mu_1(s_1 - s_2')}{1 - s_2'} dF(\theta) - M\int_{\tilde{\theta}}^{\tilde{\theta}'} \theta\mu_1 \left(1 - \frac{s_1 - s_2'}{1 - s_2'}\right) dF(\theta) \\
&= Q\mu_1(s_1 - s_2') \left[\frac{\int_{\underline{\theta}}^{\tilde{\theta}'} \theta dF(\theta)}{\int_{\underline{\theta}}^{\tilde{\theta}'} dF(\theta)} - \frac{\int_{\tilde{\theta}}^{\tilde{\theta}'} \theta dF(\theta)}{\int_{\tilde{\theta}}^{\tilde{\theta}'} dF(\theta)} \right] \\
&< 0
\end{aligned}$$

so that welfare decreases. It is possible (though not necessary) that average price decreases as well. For instance, if the distribution of the taste parameter is assumed to be uniform, then it is easy to show that $\Delta p_m = -\frac{Q}{M}(s_1 - s_2')(1 - s_1) < 0$. For an example where average price increases, consider the case where the density of θ is given by $f(\theta) = \frac{\theta^{0.1}(1-\theta)^{0.1}}{\int_0^1 \theta^{0.1}(1-\theta)^{0.1} d\theta}$, $\frac{Q}{M} = 0.99$, $s_1 = 0.5$, and $s_2' = 0.4$.

List of Tables

Table 1 Summary statistics

Variable	Mean	Median	Std. dev.	Min.	Max
Price (1969 Franc/hl)	108.9	92.5	87.1	8.4	1,784.8
Production (million liters)	70.6	20.4	156.7	0.0	1,538.3
Acreage (kha)	18.3	7.5	30.1	0.0	199.0
Yield (hl/ha)	29.2	27.0	14.8	0.5	258.0
Share of eligible acreage (1 layer)	0.064	0.000	0.198	0.000	1.000
Share of eligible acreage (5 layers)	0.004	0.000	0.028	0.000	0.336

Table 2 Trends regressions

	Price trend				Output trend	
	1927–1956		1907–1936		1927–1956	
	(1)	(2)	(3)	(4)	(5)	(6)
AOC share	78.90	75.92	-8.33	-5.77	3.87	5.21
	(28.27)	(26.38)	(6.53)	(6.97)	(11.84)	(8.27)
Vignoble FE	No	Yes	No	Yes	No	Yes
Observations	72	72	72	72	72	72

Note: The sample is limited to departments with enough information to compute price and output trends over the two periods 1907–1936 and 1927–1956. Heteroskedasticity-robust standard errors are reported in brackets. The *vignoble* control includes 16 different wine regions.

Table 3 Effect of the AOC eligible share on the real price of wine, 1907–1969

	Dep. var.: log average real price of wine									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
AOC Share1	0.416	0.463	0.453	0.448	0.441	0.380	0.437	0.427	0.418	0.412
	(0.044)	(0.045)	(0.045)	(0.046)	(0.045)	(0.045)	(0.046)	(0.046)	(0.047)	(0.046)
	[0.087]	[0.079]	[0.078]	[0.082]	[0.079]	[0.094]	[0.084]	[0.083]	[0.087]	[0.085]
	{0.108}	{0.126}	{0.127}	{0.128}	{0.130}	{0.120}	{0.136}	{0.136}	{0.139}	{0.140}
AOC Share5	–	–	–	–	–	1.253	1.407	1.407	1.497	1.491
						(0.213)	(0.250)	(0.252)	(0.269)	(0.265)
						[0.354]	[0.380]	[0.378]	[0.398]	[0.372]
						{0.494}	{0.567}	{0.547}	{0.592}	{0.573}
log(Production)	-0.041	-0.040	–	–	–	-0.038	-0.042	–	–	–
	(0.011)	(0.012)				(0.011)	(0.012)			
	[0.012]	[0.016]				[0.012]	[0.016]			
	{0.012}	{0.018}				{0.011}	{0.019}			
log(Production$_{-1}$)	–	–	-0.026	–	–	–	–	-0.027	–	–
			(0.012)					(0.012)		
			[0.017]					[0.017]		
			{0.016}					{0.017}		
log(Production)×*vignoble*	No	No	No	Yes	No	No	No	No	Yes	No
log(Production$_{-1}$)×*vignoble*	No	No	No	No	Yes	No	No	No	No	Yes
Year×*vignoble* FE	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Observations	4,572	4,572	4,483	4,572	4,483	4,572	4,572	4,483	4,572	4,483

Note: All regressions include year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years. Price data is missing for the year 1949.

Table 4 Effect of the AOC eligible share on the real price of wine, 1911–1960

	Dep. var.: log average real price of wine									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
AOC Share1	0.301	0.336	0.332	0.331	0.329	0.269	0.311	0.307	0.304	0.303
	(0.051)	(0.053)	(0.053)	(0.053)	(0.052)	(0.052)	(0.054)	(0.054)	(0.054)	(0.054)
	[0.074]	[0.074]	[0.075]	[0.075]	[0.075]	[0.081]	[0.079]	[0.079]	[0.080]	[0.080]
	{0.093}	{0.116}	{0.117}	{0.120}	{0.121}	{0.105}	{0.126}	{0.127}	{0.132}	{0.132}
AOC Share5	–	–	–	–	–	1.222	1.498	1.480	1.616	1.533
						(0.280)	(0.350)	(0.352)	(0.361)	(0.362)
						[0.231]	[0.259]	[0.254]	[0.281]	[0.245]
						{0.290}	{0.324}	{0.311}	{0.374}	{0.337}
log(Production)	-0.069	-0.066	–	–	–	-0.067	-0.068	–	–	–
	(0.011)	(0.014)				(0.011)	(0.014)			
	[0.014]	[0.016]				[0.014]	[0.016]			
	{0.014}	{0.017}				{0.013}	{0.018}			
log(Production$_{-1}$)	–	–	-0.044	–	–	–	–	-0.046	–	–
			(0.013)					(0.014)		
			[0.018]					[0.018]		
			{0.016}					{0.016}		
log(Production)×*vignoble*	No	No	No	Yes	No	No	No	No	Yes	No
log(Production$_{-1}$)×*vignoble*	No	No	No	No	Yes	No	No	No	No	Yes
Year×*vignoble* FE	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Observations	3,644	3,644	3,627	3,644	3,627	3,644	3,644	3,627	3,644	3,627

Note: All regressions include year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years. Price data is missing for the year 1949.

Table 5 Effect of the AOC eligible share on the real price of wine, 1921–1950

	Dep. var.: log average real price of wine									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
AOC Share1	0.245	0.240	0.236	0.219	0.221	0.207	0.209	0.205	0.188	0.190
	(0.067)	(0.067)	(0.066)	(0.065)	(0.063)	(0.066)	(0.067)	(0.065)	(0.065)	(0.063)
	[0.065]	[0.070]	[0.071]	[0.070]	[0.070]	[0.069]	[0.074]	[0.074]	[0.074]	[0.073]
	{0.057}	{0.074}	{0.076}	{0.076}	{0.079}	{0.066}	{0.085}	{0.086}	{0.088}	{0.089}
AOC Share5	–	–	–	–	–	1.982	2.506	2.493	2.552	2.501
						(0.527)	(0.657)	(0.666)	(0.637)	(0.667)
						[0.298]	[0.378]	[0.389]	[0.383]	[0.390]
						{0.185}	{0.104}	{0.111}	{0.083}	{0.105}
log(Production)	-0.069	-0.074	–	–	–	-0.067	-0.075	–	–	–
	(0.017)	(0.021)				(0.017)	(0.021)			
	[0.019]	[0.026]				[0.019]	[0.026]			
	{0.021}	{0.037}				{0.020}	{0.037}			
log(Production$_{-1}$)	–	–	-0.036	–	–	–	–	-0.037	–	–
			(0.022)					(0.022)		
			[0.026]					[0.026]		
			{0.034}					{0.034}		
log(Production)×*vignoble*	No	No	No	Yes	No	No	No	No	Yes	No
log(Production$_{-1}$)×*vignoble*	No	No	No	No	Yes	No	No	No	No	Yes
Year×*vignoble* FE	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Observations	2,172	2,172	2,164	2,172	2,164	2,172	2,172	2,164	2,172	2,164

Note: All regressions include year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years. Price data is missing for the year 1949.

Table 6 Effect of the later AOC eligible share on the real price of wine, pre-treatment periods

	Dep. var.: log average real price of wine								
	1907–1936					1907–1926		1917–1936	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
AOC Share1	-0.065	-0.026	-0.042	-0.025	-0.043	0.064	0.035	-0.118	-0.124
	(0.057)	(0.054)	(0.055)	(0.054)	(0.056)	(0.062)	(0.065)	(0.072)	(0.071)
	[0.067]	[0.085]	[0.086]	[0.084]	[0.082]	[0.081]	[0.079]	[0.062]	[0.065]
	{0.055}	{0.080}	{0.077}	{0.080}	{0.072}	{0.047}	{0.048}	{0.083}	{0.089}
log(Production)	-0.054	-0.062							
	(0.016)	(0.016)							
	[0.014]	[0.019]	–	–	–	–	–	–	–
	{0.012}	{0.020}							
log(Production$_{-1}$)	–	–	-0.045	–	–	–	–	–	–
			(0.014)						
			[0.021]						
			{0.026}						
log(Production)×*vignoble*	No	No	No	Yes	No	Yes	No	Yes	No
log(Production$_{-1}$)×*vignoble*	No	No	No	No	Yes	No	Yes	No	Yes
Year×*vignoble* FE	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,249	2,249	2,138	2,249	2,138	1,493	1,389	1,498	1,482

Note: All regressions include year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years.

Table 7 Effect of the AOC eligible share on the real price of wine, 1938–1969

	Dep. var.: log average real price of wine				
	(1)	(2)	(3)	(4)	(5)
AOC Share1	0.186 (0.092) [0.047] {0.057}	0.315 (0.135) [0.142] {0.048}	0.320 (0.136) [0.133] {0.038}	0.327 (0.143) [0.158] {0.036}	0.319 (0.141) [0.143] {0.033}
log(Production)	0.005 (0.019) [0.022] {0.030}	0.016 (0.022) [0.027] {0.045}	–	–	–
log(Production$_{-1}$)	–	–	0.041 (0.021) [0.026] {0.031}	–	–
log(Production)×*vignoble*	No	No	No	Yes	No
log(Production$_{-1}$)×*vignoble*	No	No	No	No	Yes
Year×*vignoble* FE	No	Yes	Yes	Yes	Yes
Observations	2,267	2,267	2,261	2,267	2,261

Note: All regressions include year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years. Price data is missing for the year 1949.

Table 8 Effect of the AOC eligible share on the real price of wine, removing non-AOC departments

	Dep. var.: log average real price of wine									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
AOC Share1	0.378	0.456	0.448	0.433	0.437	0.373	0.386	0.380	0.356	0.360
	(0.055)	(0.057)	(0.057)	(0.058)	(0.057)	(0.056)	(0.060)	(0.059)	(0.062)	(0.060)
	[0.116]	[0.127]	[0.124]	[0.126]	[0.123]	[0.127]	[0.160]	[0.155]	[0.160]	[0.156]
	{0.145}	{0.156}	{0.146}	{0.166}	{0.157}	{0.158}	{0.203}	{0.191}	{0.216}	{0.205}
AOC Share5	–	–	–	–	–	1.033	1.538	1.504	1.587	1.571
						(0.208)	(0.253)	(0.258)	(0.277)	(0.270)
						[0.434]	[0.409]	[0.404]	[0.410]	[0.382]
						{0.609}	{0.586}	{0.556}	{0.579}	{0.561}
log(Production)	-0.120	-0.087	–	–	–	-0.109	-0.100	–	–	–
	(0.028)	(0.034)				(0.027)	(0.034)			
	[0.026]	[0.040]				[0.024]	[0.044]			
	{0.026}	{0.035}				{0.026}	{0.041}			
log(Production$_{-1}$)	–	–	-0.041	–	–	–	–	-0.053	–	–
			(0.027)					(0.028)		
			[0.033]					[0.037]		
			{0.028}					{0.036}		
log(Production)×*vignoble*	No	No	No	Yes	No	No	No	No	Yes	No
log(Production$_{-1}$)×*vignoble*	No	No	No	No	Yes	No	No	No	No	Yes
Year×*vignoble* FE	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Observations	1,922	1,922	1,891	1,922	1,891	1,922	1,922	1,891	1,922	1,891

Note: The period is 1907–1969. All regressions include year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is balanced.

Table 9 Effect of AOC recognition on acreage and yield

	Dep. var.:						
	log acreage			log yield			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1921–1950	1911–1960	1907–1969	1928–1945	1921–1950	1911–1960	1907–1969
AOC Share1	0.126	0.316	0.396	-0.005	0.028	-0.026	0.073
	(0.029)	(0.040)	(0.041)	(0.078)	(0.057)	(0.046)	(0.040)
	[0.087]	[0.148]	[0.118]	[0.099]	[0.085]	[0.090]	[0.086]
	{0.058}	{0.152}	{0.142}	{0.076}	{0.072}	{0.097}	{0.102}
Observations	2,247	3,745	4,652	1,346	2,247	3,742	4,647

Note: All regressions include year by *vignoble* FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years.

Table 10 Effect of the AOC eligible share on the real price of wine, excluding departments with plausible reshuffling

	Dep. var.: log average real price of wine							
	1921–1950				1907–1969			
	Selected departments		All departments		Selected departments		All departments	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AOC Share1	0.322	0.322	0.219	0.221	0.671	0.652	0.448	0.441
	(0.103)	(0.101)	(0.065)	(0.063)	(0.062)	(0.062)	(0.046)	(0.045)
	[0.095]	[0.093]	[0.070]	[0.070]	[0.074]	[0.077]	[0.082]	[0.079]
	{0.051}	{0.053}	{0.076}	{0.079}	{0.101}	{0.110}	{0.128}	{0.130}
log(Production)×*vignoble*	Yes	No	Yes	No	Yes	No	Yes	No
log(Production$_{-1}$)×*vignoble*	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1,477	1,469	2,172	2,164	3,336	3,267	4,572	4,483

Note: All regressions include *vignoble*-by-year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years.

Table 11 Tests of the linearity assumption

	Dep. var.: log average real price of wine							
	quadratic				categorical			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AOC Share1 (β_1)	0.280 (0.107) [0.246] {0.326}	0.248 (0.106) [0.248] {0.335}	0.354 (0.111) [0.242] {0.328}	0.274 (0.108) [0.251] {0.345}	–	–	–	–
(AOC Share1)^2 (β_2)	0.209 (0.124) [0.289] {0.410}	0.233 (0.125) [0.290] {0.411}	0.108 (0.130) [0.290] {0.414}	0.191 (0.128) [0.295] {0.428}	–	–	–	–
$I_{0<s\leq 0.05}$ (β_1)	–	–	–	–	0.082 (0.024) [0.053] {0.040}	0.082 (0.024) [0.051] {0.042}	0.079 (0.024) [0.051] {0.039}	0.075 (0.024) [0.051] {0.044}
$I_{0.05<s\leq 0.40}$ (β_2)	–	–	–	–	0.112 (0.029) [0.054] {0.061}	0.103 (0.028) [0.055] {0.064}	0.128 (0.030) [0.052] {0.063}	0.105 (0.028) [0.054] {0.068}
$I_{0.40<s\leq 1}$ (β_3)	–	–	–	–	0.405 (0.038) [0.074] {0.111}	0.398 (0.038) [0.073] {0.111}	0.395 (0.039) [0.072] {0.108}	0.389 (0.038) [0.071] {0.111}
log(Prod)	Yes	No	No	No	Yes	No	No	No
log(Prod$_{-1}$)	No	Yes	No	No	No	Yes	No	No
log(Prod)×*vignoble*	No	No	Yes	No	No	No	Yes	No
log(Prod$_{-1}$)×*vignoble*	No	No	No	Yes	No	No	No	Yes
Joint significance test: $\beta_1 = \beta_2 = 0$	34.693	33.502	31.868	31.219	–	–	–	–
Linearity test: $\frac{\beta_1}{\mu_1} = \frac{\beta_2}{\mu_2}$, $\frac{\beta_1}{\mu_1} = \frac{\beta_3}{\mu_3}$	–	–	–	–	2.640	3.013	2.097	2.266
Observations	4,572	4,483	4,572	4,483	4,572	4,483	4,572	4,483

Note: All regressions include *vignoble*-by-year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years. Price data is missing for the year 1949. The Wald test of linearity in the categorical model uses the mean AOC shares across observations within each category (μ_i). The two specification tests use the department-clustered covariance matrix. The critical value for both tests is $\chi^2_{95\%}(2) = 5.991$.

Table 12 Other robustness checks, selected years

	Dep. var.: log average real price of wine											
	w/o 1940–45				w/o 1945–47				w/o 1941–1947			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
AOC Share1	0.516	0.498	0.495	0.476	0.435	0.427	0.405	0.398	0.488	0.476	0.463	0.450
	(0.048)	(0.048)	(0.049)	(0.048)	(0.048)	(0.047)	(0.049)	(0.048)	(0.049)	(0.049)	(0.050)	(0.050)
	[0.101]	[0.098]	[0.106]	[0.103]	[0.084]	[0.082]	[0.090]	[0.087]	[0.103]	[0.100]	[0.108]	[0.105]
	{0.164}	{0.164}	{0.173}	{0.174}	{0.139}	{0.139}	{0.150}	{0.150}	{0.167}	{0.168}	{0.178}	{0.179}
AOC Share5	–	–	0.959	0.994	–	–	1.464	1.495	–	–	1.190	1.237
			(0.239)	(0.222)			(0.264)	(0.250)			(0.259)	(0.237)
			[0.505]	[0.481]			[0.449]	[0.440]			[0.485]	[0.477]
			{0.668}	{0.641}			{0.686}	{0.682}			{0.719}	{0.709}
log(Prod)×*vignoble*	No	No	Yes	No	No	No	Yes	No	No	No	Yes	No
log(Prod$_{-1}$)×*vignoble*	No	No	No	Yes	No	No	No	Yes	No	No	No	Yes
Observations	4,133	4,049	4,133	4,049	4,344	4,258	4,344	4,258	4,054	3,970	4,054	3,970

Note: All regressions include *vignoble*-by-year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years. Price data is missing for the year 1949.

Table 13 Other robustness checks, selected departments

	Dep. var.: log average real price of wine							
	w/o Champagne dpts.				w/o Gironde			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AOC Share1	0.331	0.322	0.283	0.275	0.484	0.473	0.468	0.457
	(0.038)	(0.037)	(0.038)	(0.038)	(0.048)	(0.048)	(0.049)	(0.048)
	[0.074]	[0.074]	[0.081]	[0.081]	[0.084]	[0.082]	[0.084]	[0.082]
	{0.099}	{0.099}	{0.107}	{0.107}	{0.128}	{0.132}	{0.130}	{0.134}
AOC Share5	–	–	1.652	1.646	–	–	1.764	1.741
			(0.263)	(0.261)			(0.287)	(0.286)
			[0.347]	[0.321]			[0.292]	[0.265]
			{0.468}	{0.448}			{0.318}	{0.318}
log(Prod)×*vignoble*	Yes	No	Yes	No	Yes	No	Yes	No
log(Prod$_{-1}$)×*vignoble*	No	Yes	No	Yes	No	Yes	No	Yes
Observations	4,324	4,239	4,324	4,239	4,510	4,422	4,510	4,422

Note: All regressions include *vignoble*-by-year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years. Price data is missing for the year 1949.

Table 14 Time-varying effect of the AOC eligible share on the real price of wine

	Dep. var.: log average real price of wine				
	(1)	(2)	(3)	(4)	(5)
AOC Share1	0.045	0.035	0.034	0.023	0.027
	(0.064)	(0.061)	(0.061)	(0.058)	(0.059)
	[0.058]	[0.053]	[0.053]	[0.055]	[0.054]
	{0.057}	{0.054}	{0.055}	{0.057}	{0.059}
AOC Share1×(t-1937)	0.023	0.026	0.025	0.026	0.025
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
	[0.006]	[0.005]	[0.004]	[0.005]	[0.005]
	{0.008}	{0.009}	{0.009}	{0.009}	{0.009}
log(Production)	-0.053	-0.058			
	(0.011)	(0.012)	–	–	–
	[0.014]	[0.015]			
	{0.013}	{0.017}			
log(Production$_{-1}$)	–	–	-0.042	–	–
			(0.012)		
			[0.017]		
			{0.016}		
log(Production)×*vignoble*	No	No	No	Yes	No
log(Production$_{-1}$)×*vignoble*	No	No	No	No	Yes
Year×*vignoble* FE	No	Yes	Yes	Yes	Yes
Observations	4,572	4,572	4,483	4,572	4,483

Note: All regressions include year FE. Standard errors in () allow for spatial correlation up to the fourth neighboring departments and assume no time correlation. Standard errors in [] are clustered at the level of the department. Standard errors in {} are clustered at the level of the *vignoble*. The panel is unbalanced and includes all departments for which price data is available for at least half of the sample years. Price data is missing for the year 1949.