What Happens in Paris, Does Not Stay in Paris: Trade Fairs and Search and Matching Frictions*

Gábor Békés† Mátyás Molnár‡ Claudia Steinwender§

March 26, 2026

Abstract

Search and matching frictions prevent firms from forming international trade linkages. Firms commonly use trade fairs to overcome these frictions. However, despite often being subsidized by export promotion agencies, we lack evidence on how trade fairs facilitate link formation. We exploit a unique feature of Hungarian firms' participation in the 1900 Paris World Exhibition, where a trial exhibition revealed firms' ex-ante export potential category to develop a novel bounding strategy that compares treated firms to control groups from "above" and "below" in export potential. To implement our empirical strategy, we constructed a novel panel dataset of approximately 3,700 Hungarian manufacturing firms for the 1896–1906 period by digitizing, parsing and linking over 12,000 records across seven historical sources, including exhibition catalogs, government surveys, commercial directories, and official gazettes. We find that participation increases export probability by 7.6–13.3% and employment by 16.2–23.2% over eight years. Effects are larger when firms face fewer domestic competitors and more potential international buyers. This highlights both matching benefits and congestion effects when search and matching frictions are reduced.

Keywords: buyer-supplier links, export promotion, trade fairs, search frictions, industrial policy, economic history

JEL Codes: F14, N73, O14, L25

*This research is funded by the Austrian Science Fund (FWF) 10.55776/PIN3964525 and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 564732427.

†Central European University, Vienna, CEPR, ELTE KRTK Institute of Economics. bekesg@ceu.edu

‡Central European University, Vienna. molnar_matyas@phd.ceu.edu

§Ludwig-Maximilians-Universität, Munich, CEPR, CEP, CESifo. claudia.steinwender@econ.lmu.de

1 Introduction

Exporting and importing require firms to search for and connect with foreign buyers and suppliers. These firm-to-firm links are often hindered by substantial search and matching frictions. For example, a McKinsey 2021 report suggests that identifying a single suitable supplier can be costly and take several months (McKinsey & Company, 2021). Understanding the nature and magnitude of these frictions, as well as how they can be alleviated, is central to our understanding of international trade, the organization of global production networks, and the propagation of shocks across borders.

We contribute to this literature by studying one of the most common channel firms use to facilitate link formation: trade fairs. Indeed, export promotion agencies in at least 103 countries provide support to exporting firms (Lederman et al., 2010), with trade fair participation among the most common interventions. Unlike interventions that provide individual firms with targeted access to new buyers or suppliers (e.g., Atkin et al. 2017; Cai et al. 2024; Buus et al. 2025; Munch and Schaur 2018; Carballo and Volpe Martincus 2008), trade fairs bring together large numbers of firms simultaneously, exposing participants not only to potential partners but also to competitors, potentially generating congestion effects. At the same time, in contrast to transport and communication infrastructure improvements that broadly reduce effective distance between all buyers and sellers (e.g., Jensen 2007; Jensen and Miller 2018; Allen 2014; Bernard et al. 2019), trade fairs limit participation to a selected subset of firms, allowing for a more focused environment for link formation.

We overcome the selection bias that complicates the evaluation of trade fair participation by exploiting a unique historical experiment: the Hungarian government's selection process for the 1900 Paris World Exhibition. Our setting offers three advantages. First, the government used a preceding national exhibition as a "trial run," where an independent jury awarded firms based on their export potential. Second, constraints such as limited pavilion space, the desire for balanced industrial and regional representation, and transport costs of exhibits prevented the government from sending only the firms with the highest export potential. Third, the limited communication infrastructure of the era—costly travel and reliance on postal mail—meant firms had few alternatives to in-person meetings for establishing international contacts, allowing us to isolate the exhibition's impact from other correlated, private search channels that abound in modern settings.

The 1900 Paris *Exposition Universelle*, which attracted 48 million visitors and featured 76,000 exhibitors from 56 countries, was one of the largest international commercial gatherings in history. It served as a showcase for recent innovations, including the diesel engine, moving sidewalk, motion picture, and passenger trolleybus. In 1900, Hungary was part of the Austro-Hungarian Monarchy (AHM), a polity with central foreign policy and customs union, but devolved finances, economic and industrial policies.

Recognizing world fairs as an opportunity to promote Hungarian industry, the government began preparations 13 years in advance. A central element was the *1896 Budapest Millennial Exhi-*

bition, organized as a "trial run" to identify firms with the greatest export potential. The 1896 fair provided broad coverage of Hungarian firms, and an independent jury assessed each exhibit's export potential, awarding four prize categories. Based on this, a separate selection committee—distinct from the jury—was tasked with choosing firms for Paris. While the primary goal was to promote exports by selecting high-potential firms, internal records show that the committee also sought balanced representation across industries and regions and was further constrained by limited exhibition space and the high transport costs of exhibits. As a result, not all applicants could attend, and some high-potential firms were excluded in favor of others chosen to meet diversity and logistical requirements.

We exploit two features of this setting for identification. First, we use the prizes awarded in the trial exhibition to obtain an objective, ex-ante, binned measure of firms' export potential. Second, we leverage the government's constraints, which prevented perfect sorting and resulted in treated and untreated firms with similar export potential bins.

To address selection bias and estimate a causal treatment effect from attending the 1900 Paris Exhibition, we develop a novel estimation strategy that constructs "causal bands" for the treatment effect by creating two distinct control groups for treated firms, yielding upper and lower bound estimates. The upper bound compares selected firms to untreated firms *within* the same prize category, which may remain upward biased due to residual selection. The lower bound "stacks the deck" against finding a treatment effect by using a control group from the prize category immediately above the treated firms, providing a conservative estimate. This strategy also lends itself to a validation test: the upper bound estimate should exceed the lower bound, thereby confirming the assumptions about the selection process. Our estimation strategy improves upon alternative approaches such as matching estimators by exploiting additional information about whether potential controls are "close from above" or "close from below" in export potential rather than relying solely on "close" matches.

We build a novel panel dataset of Hungarian firms from seven independent historical sources, including exhibition catalogs, government surveys, commercial directories, and official gazettes for the period between 1896 to 1906. We digitized and systematically processed over 12,000 firm records using automated text extraction, geographic standardization algorithms, and multi-criteria entity resolution to construct the first comprehensive panel of Hungarian industrial firms during this period. Our dataset contains information on export status, employment, and other firm-specific variables for approximately 3,700 Hungarian manufacturing firms tracked from 1896 to 1905, although not all variables are available in every year.

We validate our estimation strategy by examining pre-treatment characteristics for both the upper and lower bound approaches. As expected, in the upper bound estimation, treated firms slightly outperform untreated firms in several baseline characteristics. In contrast, in the lower bound estimation, there is no significant difference between treated and untreated firms; if anything, treated firms perform slightly worse at baseline.

For our main outcome of interest—exporting—we have two pre-exhibition observations, which allow us to test for pre-trends. Consistent with expectations, the upper bound estimation shows a pre-trend, suggesting that it may be upward biased. In contrast, the lower bound estimation exhibits no pre-trend, supporting its interpretation as a more conservative estimate. Overall, these results suggest that the true causal effect likely lies closer to the lower bound, though within the estimated range.

We find that exhibition participation in Paris increased the export probability 4 to 5 years after the exhibition. For the broader export capability status, the causal band corresponds to a 7.6 to 13.3 percentage point higher probability in 1904. For the narrower definition of extra AHM export status, the band reflects a 4.2 to 10 percentage point higher probability in 1905. Beyond exports, we document substantial increases in employment among Paris exhibitors, with estimated gains of 16.3–23% over an eight-year period, equivalent to roughly 1.8–2.6 percentage points of annual growth. However, these employment gains are not accompanied by increases in mechanization.

At the same time, this average effect masks substantial heterogeneity that sheds light on search and matching frictions. Moving from the 25th to the 75th percentile of domestic competitors reduces the treatment effect by two-thirds, consistent with congestion effects. Conversely, moving across the same range of potential international buyers increases it by nearly half, consistent with the benefits of market thickness.

Our work contributes to four strands of literature. First, we relate to a growing literature studying firm-to-firm relationships. The empirical literature has focused on showing that exogenous reductions in search and matching frictions cause more matches. Part of the literature has used quasi-experiments that reduce transportation and communication cost to show evidence of search and matching frictions (e.g., Jensen (2007); Jensen and Miller (2018); Allen (2014); Bernard et al. (2019)). Others employed RCTs to measure the effect of giving firms access to new suppliers (Cai et al., 2024) or buyers (Atkin et al., 2017; Bergquist et al., 2024).

A complementary theoretical literature incorporates search and matching frictions or fixed costs of link formation to explain common patterns of supplier-buyer relationships (e.g.,Bernard et al. 2019; Demir et al. 2024; Arkolakis et al. 2025; Eaton et al. 2022a), Dhyne et al. 2023; Lim 2017; Bernard et al. 2018, Eaton et al. 2022b; Miyauchi 2024). For international trade. Lenoir et al. (2023) structurally estimated search frictions using French firm-level data, finding them as important as traditional iceberg costs in impeding trade. Krolikowski and McCallum (2021) embedded goods-market frictions in a general equilibrium framework and show that matching frictions alter the gains from trade by creating an endogenous fraction of unmatched exporters.

We contribute to this literature by studying trade fairs—the most common mechanism firms use to form new links—which, unlike targeted interventions, simultaneously expose firms to potential partners and competitors, and unlike infrastructure improvements, operate in a selective, focused environment for matching. Our findings on the heterogeneous effects of trade fairs across product categories with different levels of competition and buyer concentration speak directly to these

mechanisms, providing causal evidence on how search and matching frictions operate when firms are brought together in a concentrated environment.

Second, we contribute to the literature evaluating export promotion programs (e.g., Munch and Schaur 2018; Martincus and Carballo 2010; Carballo and Volpe Martincus 2008; Broocks and Van Biesebroeck 2017; Munch et al. forthcoming; Van Biesebroeck et al. 2015; Hiller 2012; Görg et al. 2008; Schminke and Van Biesebroeck 2013). While early studies often relied on difference-in-differences matching based on observable firm characteristics, more recent work exploits quasi-random variation in the selection processes of export promotion agencies to identify causal effects. We extend this literature in three ways. First, we introduce a novel estimation approach that establishes "causal bands" for treatment effects, improving upon prior matching estimators by distinguishing between comparisons "close from above" and "close from below." Second, whereas export promotion agencies sometimes subsidize trade fair participation as part of a broader package of services, we focus exclusively on trade fairs as a standalone intervention. Third, our setting allows us to examine not only the average effect on treated firms but also the role of competition among participants, highlighting how exposure to competitors can generate congestion effects.

Third, our work contributes to the growing body of quantitative research on how industrial policy influenced development during the first era of globalization. While the historical literature has documented the prevalence of industrial policy (see, e.g., the survey in (Juhász and Steinwender, 2024)), quantitative evaluations remain scarce. Notable exceptions include studies on temporary protection (Juhász, 2018; Hanlon, 2020), infrastructure improvements (Steinwender, 2018), intellectual property protection (Alfaro et al., 2022), colonialism (Romero, 2023), and tariffs (Juhász, 2018). In contrast, the role of trade fairs in shaping firm outcomes remains unexplored. Prior research has instead used trade fair exhibitor data as an alternative measure of innovation to examine questions related to intellectual property protection (Moser, 2005, 2012). We study a large-scale, centrally planned effort that transformed Hungary from a minor participant in 1878 into the fifth-largest national exhibitor at the 1900 fair.

Finally, our estimation approach builds on the literature using bounds to address selection problems in program evaluation, deriving credible treatment effect bounds under weaker assumptions (e.g., Manski 2003; Lee 2009; Manski and Pepper 2000; Tamer 2010; Ho and Rosen 2017; Blundell et al. 2007). Our setting shares key features with this literature: we face a selection problem where exhibiting firms likely differ from non-exhibitors in unobservable ways. Like Manski and Pepper (2000), we exploit an ex-ante ranking variable—the Budapest exhibition awards—that partially predicts treatment assignment. However, our approach innovates in how we construct bounds. Rather than using trimming procedures as in Lee (2009), or comparing across levels of the monotone treatment selection pattern as in Manski and Pepper (2000), we exploit the discrete nature of our award categories to construct two distinct control groups (same vs. lower-tier awards). This approach offers several advantages: it provides clear economic intuition by comparing treated firms to those of similar versus superior export potential and incorporates a built-in validation test, as

the upper bound should exceed the lower bound. We compare our bounds to the Manski bounds in Section 6.1.

Section 2 provides the historical context. Section 3 describes the construction of a panel dataset of Hungarian firms by digitizing and linking multiple historical sources. Section 4 outlines a novel empirical strategy for constructing upper and lower bounds on the treatment effect. Section 5 details the application of this empirical strategy to the setting, and Section 6 presents the results. We then look at heterogeneity of treatment effects across product subcategories to understand the role of competition and market thickness in Section 7. We conclude in Section 8 and relate results to contemporary debates.

2 Hungarian Firms in the 1900 Paris World Exhibition

The 1900 Paris *Exposition Universelle* was the largest world's fair to date, celebrating the achievements of the 19th century and showcasing visions for the 20th. Running from April to November 1900, it attracted over 48 million visitors, roughly equivalent to France's entire population. The exhibition featured 76,000 exhibitors from 56 countries displaying innovations ranging from moving sidewalks (*trottoirs roulants*) and early motion pictures to diesel engines and passenger trolleybuses. Exhibitors were organized into 18 broad categories based on the general theme of the products they exhibited, and further subdivided into 121 subcategories corresponding to narrower product groups.

The exhibition took place across multiple halls in different parts of Paris, many of them near the *Champs de Mars*. Firms in the same broad category shared a hall; for instance, one hall was dedicated to "Mines et métallurgie" (mining and metallurgy). Within each hall, exhibits were arranged by country, with Hungary typically placed alongside Austria and Belgium.

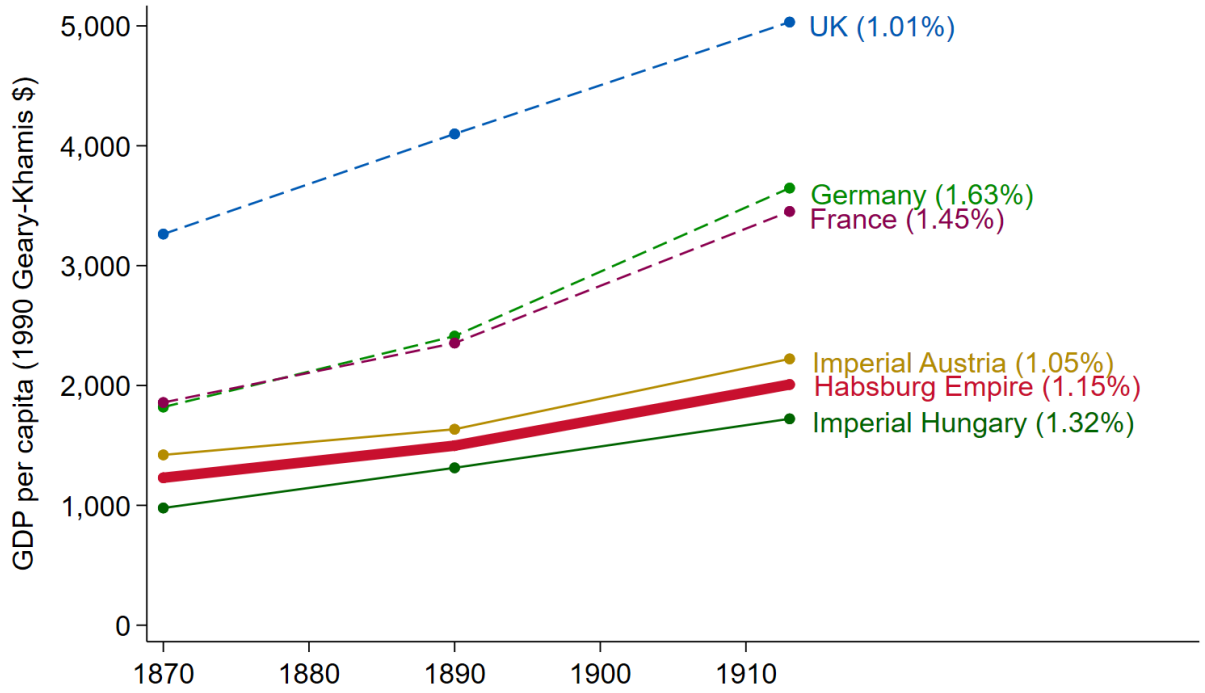
For participating countries, the exhibition served purposes well beyond cultural diplomacy and national prestige. Governments explicitly recognized its function as, in effect, a trade fair: the vast concentration of international visitors, including merchants, industrialists, and potential customers, created unprecedented opportunities for establishing business connections (Tamilia, 2007). Firms showcased their products and innovations, demonstrated product capabilities, negotiated terms with prospective buyers, and used the occasion for networking through meetings, lunches, dinners, and organized events. Contemporary accounts describe the exhibition grounds as bustling with commercial negotiations alongside the display of innovations.

The exhibition's commercial importance was amplified by the communication constraints of the era. International business relationships typically required either costly travel or slow correspondence by mail. The telegraph, while revolutionary for time-sensitive information, was most effective for trade in homogeneous products that could be codified, but limited in facilitating trade in goods that required physical inspection (Juhász and Steinwender, 2018). The exhibition thus offered a rare opportunity for face-to-face meetings that could build trust and enable detailed product demonstrations.

Between 1867 and 1918, Hungary was part of the Austro-Hungarian Monarchy (AHM), a dual state in which the Austrian Emperor also served as King of Hungary. The empire maintained a common foreign policy and customs union, but Hungary had its own parliament, finances, and economic and industrial policies.

The AHM was a late industrializer compared to Western European powers such as Great Britain, France, or Germany, resulting in substantially lower per capita incomes between 1870 and 1913, as shown in Figure 1. However, within the monarchy, Hungary grew faster than Austria, achieving an annual GDP per capita growth rate of 1.32%, closer to France (1.45%) than Austria (1.05%) (Schulze, 2000). This reflected Hungary's industrialization, which had begun in the late 19th century: the employment share in manufacturing, crafts, and utilities rose from 6.6% in 1869/1870 to 11.1% in 1910, though it still lagged behind Austria (21.4%) and Germany (29.4%) (Schulze, 2007). Among the most dynamic manufacturing sectors were food processing (especially flour milling) in the early phase (1870–90); by the 1890s and early twentieth century, engineering and metal-working received a major stimulus from empire-wide machinery demand and Austrian capital inflows, becoming important growth sectors (Klein et al., 2017).

Figure 1: GDP per capita, 1870–1913



Source: Own computations based on data from Schulze (2000), Table 8.

The Hungarian government perceived the Paris Exhibition as an ideal opportunity to promote Hungarian exports (Kun, 2019), especially since it marked the first time Hungary participated in a world fair as an exhibitor independent of Austria. The Hungarian Royal Minister of Commerce,

Baron Ernő Dániel, noted that "*[...] the 1900 Paris Exhibition [...] will be a meeting place for foreigners flocking there from all parts of the world, and thus a favorable opportunity for establishing commercial connections. The Hungarian Royal Commissioner, who [...] is [...] organizing the representation of the countries of the Hungarian Crown [...] will support the establishment of such commercial connections*" (*Hungarian Industrial Journal*, 1897, ed. 2).

To maximize the economic benefits of the exhibition, the Hungarian government prepared extensively and strategically. Preparations began as early as 1887, thirteen years before the Paris Exhibition, and included establishing institutions and procedures for selecting firms. A key element was the use of the 1896 Budapest Millennial Exhibition, held four years before Paris, as a trial exhibition to identify the country's most competitive products and firms (Székely, 2015).

To ensure broad participation in the trial exhibition, the Budapest exhibition was widely advertised in newspapers and posters, supplemented by outreach through industrial associations and chambers of commerce. The effort paid off: with 9,650 exhibitors from the Hungarian Kingdom proper participating, the exhibition provided near-comprehensive coverage of Hungary's export-relevant manufacturing sector.

At the trial exhibition, an independent jury assessed each exhibit's export potential and awarded four prize categories in descending order: gold, silver, bronze, and honorable mention.[1] The jury was organized into 34 product-group panels with a total of 480 members, half elected by exhibitors and half appointed, following formal regulations.[2] An exhibit received an award only upon a majority vote of the relevant panel.[3] Jury members and their firms were excluded from receiving prizes.[4]

After the trial exhibition, a separate selection committee, distinct from the 1896 award jury, was tasked with choosing firms for Paris. While the primary goal was to promote Hungarian exports by selecting firms with high export potential, internal documents show that the committee also sought balanced representation across industries and regions, aiming to present a broad cross-section of Hungarian production. In addition, exhibition space in Paris was limited and transport costs were high, so not all promising candidates could be accommodated. Within some product categories, space constraints meant that even firms with demonstrated export potential were excluded, while others were included to maintain sectoral balance.

Overall, Hungary sent approximately 3,000 exhibitors across all sectors to Paris, making it the fifth-largest national exhibitor at the fair. Beyond granting participation, the government covered all costs for selected exhibitors, including transportation of exhibits to Paris.

[1] For simplicity, we relabeled the names; the original names and descriptions of the awards are found in Appendix E.1.2.

[2] For details on jury selection, see Appendix E.1.1.

[3] For more details on criteria and justification for awards, see Appendix E.1.3 and E.1.4.

[4] Note that for this reason, we exclude jury members from our sample of firms.

3 Data

We combine seven datasets, listed in Table 1, to construct a panel dataset of Hungarian firms. All data were digitized by us from historical archival sources housed in the Hungarian National Library, the Austrian National Library, or available online at the Internet Archive.[5]

Primary sources, accessed via microfilm, in print, or as digital scans, underwent systematic digitization involving photographic capture, optical character recognition (OCR) using Google Document AI, and error correction through comparison with original sources. The resulting OCR output was processed to standardize Hungarian diacritical marks and historical orthographic conventions, extract structured data fields from semi-structured catalog entries, and correct extraction errors through manual verification.

Table 1: Overview of historical data sources

Name (translated)	**Original name**	**Year**	**Key variables**
Budapest Exhibition Catalog	Az 1896-iki Ezredéves Országos Kiállítás általános katalógusa	1896	Product group, firm characteristics, incl. export status and employment
Budapest Exhibition Brochure	1896-iki ezredéves országos kiállítás közleményei	1896	Awards from Budapest exhibition
Central Gazette of Hungarian Firms	Központi Értesítő	1896–1904	Firm exits
Paris Exhibition Catalog	Catalogue Général Officiel	1900	Treatment, product group
Hungarian Museum of Commerce Survey 1	Kiviteli czímtár	1898	Export capability and employment
Hungarian Museum of Commerce Survey 2	Kiviteli címtár	1904	Export capability and employment
Austro-Hungarian Export Compass	Export-Compass, commerzielles Jahrbuch für d. Interessenten des österreichisch-ungarischen Ausfuhr-Handels	1905	Export status

[5]Available at `https://archive.org`.

The most complex technical challenge was entity resolution across the independently collected datasets. Firm names vary in spelling, abbreviation, and format across sources, requiring a combination of matching algorithms and extensive manual checking of potential matches. We created correspondence tables across dataset pairs and consolidated them into a master identification system enabling both balanced and unbalanced panel specifications. Further details on the digitization process and entity resolution methods are provided in Appendix D.

In the following, we describe how we construct the dataset used in our baseline analysis combining the different sources.

Sample. The *Budapest Exhibition Catalog* documented 9,650 exhibitors from Hungary proper at the 1896 Budapest Millennial Exhibition (excluding Croatia-Slavonia and Bosnia-Herzegovina) across 23 exhibition groups: one historical group and 22 contemporary groups broadly corresponding to economic sectors. The distribution of exhibitors across groups is shown in Table B.4.

To focus on commercially oriented manufacturing firms, we apply three filters to the full catalog. First, we retain only contemporary exhibition groups within manufacturing, excluding groups such as Group I/a (Fine Arts), III (Education), or V (Trade, Finance, and Credit), where no for profit manufacturing firms are present. Second, within these groups, we exclude non-commercial entities (municipalities, associations, and NGOs) and consolidate duplicate entries where firms exhibited across multiple categories. Third, we restrict the sample to firms that entered the jury competition for awards, which excludes jury members (70 firms) and exhibitors that explicitly asked to be excluded from competition (72 firms).

After applying these filters, the sample comprises 3,736 for-profit manufacturing firms. Not all of these firms survive until the end of our sample period in 1904–1905. To estimate treatment effects conditional on survival, we identify firm exits between 1896 and 1904 by matching firm names to the exit reports from the *Central Gazette of Hungarian Firms*. We find that 122 firms exited during this period, bringing our baseline sample to 3,614 firms.[6]

Baseline characteristics. The *Budapest Exhibition Catalog* provides exhibitor-level information, including names, addresses, exhibited products, and exhibition group, for nearly all participants. For a subset of exhibitors, the catalog also records firm-level characteristics such as workforce size, skill composition, mechanization (installed horsepower), output value, export destinations, year of establishment, and patent holdings. In addition, we identify joint stock companies from the firm names and multi-location firms in cases when the dataset indicates other locations for factories, warehouses, or offices besides their official address.

Where firm characteristics are missing from the *Budapest Exhibition Catalog*, we supplement them with the *Hungarian Museum of Commerce Survey 1* from 1898 where available.[7]

[6] In Section 6.3, we examine the effect of the Paris exhibition on firm exits and survival.

[7] Since firms are observed over windows of varying length, we rescale log growth in employment to a common 8-year horizon by dividing the observed log growth rate by the number of years elapsed and then

Unobserved export potential: award category. We match these firms to the official jury results published in the *Budapest Exhibition Brochure* at the exhibition's close to obtain information on the awards granted to each firm. There were four award categories: gold, silver, bronze, and honorable mention.[8]

Treatment: participation in the Paris Exhibition. We identify treated firms in our baseline sample as Paris exhibitors if we successfully match it to an entity listed in the *Paris Exhibition Catalog*.

Outcomes: export capability and export status. We use two complementary measures of export activity: export capability and export status.

Export capability is measured by whether a firm appears in the *Hungarian Museum of Commerce Survey 1* covering 1898 (pre-treatment) or in the *Hungarian Museum of Commerce Survey 2* covering 1904 (post-treatment).[9] The stated objective of the survey was to compile a registry of firms that were either actively exporting or capable of exporting beyond the borders of the Hungarian Kingdom. The Museum defined export capability as encompassing "industrial plants that exceed local consumption demands, are prepared for international shipping, and maintain capacity to engage with foreign importers."[10] This included firms that "only lacked a foreign order to become export companies," as well as firms selling outside their home region (e.g., from Budapest to Slavonia, now part of Croatia, or to Pozsony, now Bratislava, Slovakia).[11]

Export status is measured by whether a firm appears in the *Austro-Hungarian Export Compass* from 1905, compiled by the Austrian Ministry of Commerce in partnership with regional chambers of commerce. This directory lists all firms that exported outside the Austro-Hungarian customs union in the reference year, organized by product category.

Export status is a stronger measure than export capability for two reasons. First, it captures actual exporting rather than the ability to export. Second, it records exports outside the Austro-Hungarian customs union, a higher bar than exports outside Hungary but still within the Monarchy.

Outcomes: employment. The *Hungarian Museum of Commerce Survey 2* provides data on employment levels for 1904 (post-treatment), albeit for only a subset of firms.

Summary statistics on all variables are presented in Table 2, separately for baseline characteristics and variables measured post-treatment. Of the Hungarian firms in our sample, 15% participated in the 1900 Paris Exhibition. At baseline, 15% of firms were already exporting in 1896, and 25%

multiplying it by 8.

[8]For details, see Appendix E.1.2.

[9]These surveys were conducted in 1894, 1897-98, 1904, 1909 and 1914. The earliest survey from 1894 only covered a limited sample of firms, and the last two editions from 1909 and 1914 register procurement sources rather than exporting firms, which is why we focused on 1897-98 and 1904.

[10]Authors' translation from Hungarian.

[11]For additional details, see Appendix D.

reported being capable of exporting in 1898. Detailed firm characteristics such as workforce size, installed horsepower, and output value are available for about 13% of firms. Among firms with available employment data, the range spans from 1 to over 17,000 workers, with a mean of 172. Five percent of firms hold at least one patent, 7% operate more than one establishment in Hungary, and 4% are organized as joint-stock companies. The average firm age is 26 years.[12]

Turning to the award categories from the Budapest Exhibition, gold was the most selective, received by only 3.7% of firms. Another 11.2% received silver, followed by a larger bronze category comprising 27.9% of firms. Nearly a third received an honorable mention, and 26.1% received no award.

Table 2: Summary Statistics

Variable	N	Mean	SD	Min	Max
Treatment					
Paris exhibitor dummy	3,614	0.15	0.35	0	1
Baseline characteristics (1896, unless indicated otherwise)					
Export status	3,614	0.15	0.36	0	1
Export capability (1898)	3,614	0.25	0.44	0	1
Change in exporting (export cap. 1898-export status 1896)	3,614	0.10	0.37	-1	1
All information available (worker, horsepower, output)	3,614	0.13	0.34	0	1
Worker information available	3,614	0.33	0.47	0	1
Number of workers	1,193	172.40	783.47	1	17,165
Share of skilled workers	992	0.76	0.39	0.006	8.67
Horsepower (HP) information available	3,614	0.18	0.38	0	1
Installed horsepower	652	224.89	997.51	1	15,700
Automation (HP/worker)	640	0.92	1.62	0.008	23.40
Output information available	3,614	0.15	0.36	0	1
Output (in HUF)	541	688,042	1.842 mil	400	14 mil
Labor productivity (output/worker, in HUF)	506	6,652	33,588	34.04	707,500
Patent dummy	3,614	0.05	0.21	0	1
Multi-location firm	3,614	0.07	0.25	0	1
Joint-stock company	3,614	0.04	0.19	0	1
Firm age	1,391	25.95	30.76	1	510
Gold medal indicator	3,614	0.037	0.190	0	1
Silver medal indicator	3,614	0.112	0.315	0	1
Bronze medal indicator	3,614	0.279	0.448	0	1
Honorable mention indicator	3,614	0.310	0.463	0	1
No award indicator	3,614	0.261	0.440	0	1
Post-treatment (1904, unless indicated otherwise)					
Export status (1905)	3,614	0.11	0.31	0	1
Export capability	3,614	0.24	0.43	0	1
Number of workers	796	192.01	740.57	2	17,000

[12]Appendix Table B.2 presents baseline characteristics separately for Paris exhibitors and non-Paris exhibitors.

Competition in Paris. To measure the competitive environment faced by Hungarian firms at the Paris Exhibition, we extract exhibitor counts by country and subcategory from the General Catalog.

4 Empirical Strategy: Upper and Lower Bounds on the Treatment Effect

The central challenge in estimating the causal effect of trade fair participation is selection: firms that exhibit are likely to differ systematically from those that do not. Let $Y_i(1)$ and $Y_i(0)$ denote the potential outcomes for firm i with and without trade fair participation, respectively. The firm-level treatment effect is

$$\tau_i = Y_i(1) - Y_i(0).$$

Since only one potential outcome is observed for each firm, realized outcomes satisfy

$$Y_i = D_i \cdot Y_i(1) + (1 - D_i) \cdot Y_i(0),$$

where D_i is an indicator for participation. A naïve comparison of participants and non-participants yields

$$\begin{aligned} E[Y(1)|D=1] \;-\; & E[Y(0)|D=0] = \\ = \; & E[Y(1)|D=1] - E[Y(0)|D=1] + E[Y(0)|D=1] - E[Y(0)|D=0] \\ = \; & \underbrace{E[\tau|D=1]}_{\text{ATT}} + \underbrace{E[Y(0)|D=1] - E[Y(0)|D=0]}_{\text{selection bias}}. \end{aligned}$$

If firms with greater export potential are more likely to participate, the selection term $E[Y(0)|D=1] - E[Y(0)|D=0]$ is positive, biasing the naïve difference upward relative to the average treatment effect on the treated (ATT).

We address this selection problem by exploiting a discrete proxy A_i for firms' unobserved export potential. Formally, let $A_i \in \mathcal{A} = \{1, 2, 3, 4, 5\}$, where $A_i = 1$ denotes the highest export potential. In our application, the role of A_i is played by the award categories from the preceding trial exhibition; we therefore refer to A_i as award categories throughout. We defer the discussion of why awards are a valid proxy for export potential to Section 5 and first establish the assumptions that A_i must satisfy for our bounding strategy.

The key idea is to use the award categories to construct two distinct control groups for treated firms, yielding an upper and a lower bound on the ATT. The upper bound compares treated firms to untreated firms within the same award category. Because residual positive selection may persist within categories, this comparison is likely upward biased. The lower bound instead compares treated firms to untreated firms from the award category immediately above, "stacking the deck"

against finding an effect by assigning treated firms a control group with higher export potential and thereby providing a conservative estimate.

If A_i were a continuous and exact measure of firms' underlying export potential rather than a coarse binned proxy, the natural solution would be to compare treated and untreated firms with identical values of A_i, as in a standard matching estimator. Because A_i is discrete, however, it is not granular enough to eliminate within-bin selection bias. Moreover, even with a more finely measured proxy, researchers may fail to find exact matches in practice, and nearest-neighbor matching introduces residual bias of unknown sign and magnitude. Our approach sidesteps these difficulties: rather than seeking exact matches, we construct control groups that bound the treatment effect from above and below. Our approach is therefore not limited to coarsely binned proxies; it applies equally to settings where a more fine-grained measure of unobserved potential is available.[13]

Our approach requires two assumptions on A_i, one for each bound. The first assumption establishes the upper bound by formalizing the idea that positive selection persists within award categories:

Assumption 1 (Upper bound)**.** *Within award categories, selected firms have weakly higher export potential than non-selected firms:*

$$E[Y_i(0) \mid D_i = 1, A_i = a] \geq E[Y_i(0) \mid D_i = 0, A_i = a], \quad \forall\, a.$$

Assumption 1 implies that within-award comparisons of treated and untreated firms are subject to residual positive selection bias, so that the resulting estimates are upward biased.

Lemma 1 (Upper bound)**.** *Define the upper bound estimate for each award group a as $\beta^a_{upper} := E[Y(1) \mid D = 1, A = a] - E[Y(0) \mid D = 0, A = a]$, and the award-specific ATT as $\tau^a := E[Y(1) \mid D = 1, A = a] - E[Y(0) \mid D = 1, A = a]$. Under Assumption 1,*

$$\beta^a_{upper} \geq \tau^a,$$

i.e., β^a_{upper} is an upper bound on the treatment effect on the treated.

Proof.

$$\begin{aligned} \beta^a_{upper} &= E[Y(1) \mid D = 1, A = a] - E[Y(0) \mid D = 0, A = a] \\ &= E[Y(1) \mid D = 1, A = a] - E[Y(0) \mid D = 1, A = a] \\ &\quad + E[Y(0) \mid D = 1, A = a] - E[Y(0) \mid D = 0, A = a] \\ &= \tau^a + \underbrace{E[Y(0) \mid D = 1, A = a] - E[Y(0) \mid D = 0, A = a]}_{\geq 0}. \end{aligned} \tag{1}$$

□

[13] We compare our approach to standard matching estimators in Section 6.1.

The second assumption establishes the lower bound. Rather than comparing treated firms to untreated firms within the same award category, we compare them to untreated firms from the category immediately above, who have weakly higher export potential:

Assumption 2 (Lower bound)**.** *Non-selected firms from each award category have weakly higher export potential than selected firms from the category below:*

$$E[Y_i(0) \mid D_i = 0, A_i = a-1] \geq E[Y_i(0) \mid D_i = 1, A_i = a], \quad \forall\, a \in \mathcal{A} \setminus \{1\}.$$

By comparing treated firms to a strictly better control group, this assumption generates estimates that are downward biased, providing a lower bound on the ATT.

Lemma 2 (Lower bound)**.** *Define the lower bound estimate for each award group a as $\beta^a_{lower} := E[Y(1) \mid D = 1, A = a] - E[Y(0) \mid D = 0, A = a-1]$. Under Assumption 2,*

$$\beta^a_{lower} \leq \tau^a,$$

i.e., β^a_{lower} is a lower bound on the treatment effect on the treated.

Proof.

$$\begin{aligned}
\beta^a_{lower} &= E[Y(1)|D=1,A=a] - E[Y(0)|D=0,A=a-1] \\
&= E[Y(1)|D=1,A=a] - E[Y(0)|D=1,A=a] \\
&\quad + E[Y(0)|D=1,A=a] - E[Y(0)|D=0,A=a-1] \\
&= \tau^a + \underbrace{E[Y(0)|D=1,A=a] - E[Y(0)|D=0,A=a-1]}_{\leq 0}
\end{aligned}$$

□

4.1 Empirical Implementation

To aggregate the award-class-specific bounds across award groups, we implement our bounding strategy via two regression specifications that differ in their fixed effects and sample composition.

The upper bound is estimated on the full sample, pooling observations across all award categories:

$$Y_i = \beta_{upper} \cdot D_i \; + \; \sum_a \gamma_a \cdot \mathbf{1}\{A_i = a\} \; + \; \epsilon_i, \tag{2}$$

where Y_i denotes the outcome of interest, D_i is the treatment indicator, and $\mathbf{1}\{A_i = a\}$ are award category fixed effects. The inclusion of award fixed effects ensures that treated and untreated firms are compared within the same award category, as required by the upper bound construction in

Lemma 1.[14]

The lower bound is estimated using the following specification:

$$Y_i = \beta_{lower} \cdot D_i \; + \; \sum_g \theta_g \cdot \mathbf{1}\{G_i = g\} \; + \; \epsilon_i, \tag{3}$$

where $G_i \in \{1, 2, 3, 4\}$ denotes comparison group membership. Each comparison group g pairs treated firms from award class $g + 1$ with untreated firms from award class g. The comparison group fixed effects thus ensure that treated firms are compared to untreated firms from the award category immediately above, as required by Lemma 2. This specification requires dropping treated firms in the top award class (who have no higher-ranked control group) and untreated firms in the bottom award class (who have no lower-ranked treatment group).

Equations (2) and (3) therefore differ in two respects: the employed fixed effects and the estimation sample.

4.2 Validation Tests

The credibility of our bounds rests on Assumptions 1 and 2, neither of which is directly testable. This parallels other research designs: in difference-in-differences, for instance, researchers routinely conduct balancing tests on pre-treatment variables to support the plausibility of their identifying assumptions. In our research design, this corresponds to two types of balancing checks by testing for differences in pre-treatment characteristics between treated and untreated firms within their respective comparison groups, i.e., within award classes for the upper bound (corresponding to testing Assumption 1) and across adjacent award classes for the lower bound (corresponding to testing Assumption 2).

Beyond these two sets of balancing tests, our setting lends itself to two additional validation tests. The first exploits the fact that Assumptions 1 and 2 jointly imply a complete ordering of unobserved export potential across subgroups defined by the interaction of award class and treatment status. We formalize this ordering in Lemma 3 and test whether it holds for observable pre-treatment characteristics using a Spearman rank correlation test.

Lemma 3 (Ordering of Unobserved Export Potential)**.** *Define rank groups $r_i = 2A_i - D_i$, so that $r_i \in \{1, 2, \ldots, 10\}$. Under Assumptions 1 and 2, unobserved export potential is weakly decreasing*

[14]OLS aggregates across award classes using implicit weights proportional to within-group treatment variation, yielding a variance-weighted average of group-specific ATTs. An alternative is to weight by the share of treated firms in each award class. Since the group-level inequalities hold for each award class individually, any positive-weighted average preserves the bounding interpretation. We prefer the OLS implementation for its flexibility in accommodating additional fixed effects and interaction terms. In Section 6.1, we show that the regression-based bounds and the share-weighted bounds are similar in our application.

in r:

$$\begin{aligned} E[Y_i(0) \mid r_i = 1] &\geq E[Y_i(0) \mid r_i = 2] \\ &\geq E[Y_i(0) \mid r_i = 3] \\ &\geq \cdots \\ &\geq E[Y_i(0) \mid r_i = 10] \end{aligned}$$

Proof. Odd-numbered inequalities (e.g., $r = 1 \geq r = 2$) follow from Assumption 1 (within-award positive selection), and even-numbered inequalities (e.g., $r = 2 \geq r = 3$) follow from Assumption 2 (across-award ordering). □

The second test exploits the implication that the upper bound estimate must exceed the lower bound estimate:

Lemma 4 (Consistency of Bounds)**.** *Under Assumptions 1 and 2,*

$$\beta_{upper} \;\geq\; \beta_{lower}.$$

Failure of this inequality in the data would reject at least one of the two assumptions. The condition is necessary but not sufficient: the inequality holding is consistent with, but does not guarantee, the validity of the assumptions.

4.3 Relation to the Partial Identification Literature

Our approach relates to a broader econometrics literature on partial identification and bounding treatment effects when point identification is not feasible. Classic contributions include Manski (2003), who provides a general framework for bounding treatment effects under weak assumptions, and Lee (2009), who develops trimming bounds for sample selection. More recent surveys of this literature include Tamer (2010) and Ho and Rosen (2017).

A foundational challenge in this literature is sample selection, where outcomes are observed only for a selected subsample. Lee (2009) addresses this in the context of job training programs, where wages are observed only for employed individuals. Since training may affect both employment and wages, simple comparisons among the employed suffer from selection bias. Lee imposes a strict monotonicity assumption on the selection into the observed sample and develops a trimming procedure that yields sharp bounds on the average treatment effect. Blundell et al. (2007) tackle a related problem when studying wage inequality over time with changing employment rates, bounding the wage distribution by combining the observed distribution for employed individuals with worst-case assumptions about the non-employed.

Our approach is most closely related to Manski's partial identification framework (Manski, 2003), which derives bounds under the weak requirement that outcomes are bounded. For binary

outcomes, as in our setting, these bounds have width 1 and therefore necessarily include zero, which is typically too coarse for inference. Two common refinements tighten these bounds. Monotone treatment selection (MTS) assumes that treated units have weakly higher counterfactual outcomes than untreated units, while monotone treatment response (MTR) assumes that treatment weakly improves outcomes for each treated unit.

An alternative refinement is to use an additional covariate believed to be monotonically related to potential outcomes as a *Monotone Instrumental Variable* (MIV) (Manski and Pepper, 2000). In their application to returns to schooling, Manski and Pepper observe that while AFQT scores do not perfectly determine educational attainment, they provide a credible ranking: individuals with higher test scores are more likely to obtain more schooling. The key insight is that comparing outcomes across levels of the ranking variable provides information about treatment effects even when the ranking does not fully determine treatment.

Our approach shares with the MIV framework the idea of exploiting an additional covariate that ranks units by their potential outcomes. However, the assumptions differ in an important respect. The MIV assumption requires that expected potential outcomes are monotonically related to the covariate across *all* firms, treated and untreated combined. Our assumptions are stronger: we require that *untreated* firms in a given award category have weakly higher potential outcomes than *treated* firms in the category below (Assumption 2), and that within award categories, treated firms have weakly higher potential outcomes than untreated firms (Assumption 1). For a formal comparison of our identifying assumptions to bounded outcomes, MTR, MTS, and MIV, see Appendix C.

The stronger assumptions deliver three advantages relative to the existing literature. First, they yield tighter bounds, which we illustrate in Section 6.1. Second, they generate two testable implications (Lemmas 3 and 4), increasing the credibility of the estimates. Third, our approach sidesteps the common difficulty of finding control observations that closely resemble treated units: rather than seeking exact matches to infer counterfactual outcomes, we construct control groups that bound the treatment effect from above and below, which may be easier to defend. Our approach is not limited to coarsely binned proxies and applies equally to settings where a more fine-grained measure of unobserved potential is available.

5 Application to the Historical Setting

The research design introduced in Section 4 requires a covariate A_i that ranks firms' unobserved export potential in accordance with Assumptions 1 and 2. In our setting, the awards from the trial exhibition (the 1896 Budapest Millennial Exhibition), which preceded the Paris Exhibition by four years, serve this role. Four features of the historical setting support the suitability of these awards. First, the awards were determined ex ante, before the Paris exhibition and even before participants were selected. Second, export potential was an explicit evaluation criterion of the

jury for awards.[15] Third, the awards were determined by an independent expert jury, separate from the subsequent selection committee for Paris. Fourth, government constraints, such as the objective of a balanced representation across industries and regions, as well as limited exhibition space and the high transport costs of exhibits, prevented perfect sorting of Paris exhibitors across award categories and resulted in treated and untreated firms with similar export potential bins.

While these institutional features suggest that awards are a suitable proxy for export potential, we now examine empirically whether Assumptions 1 and 2 are satisfied.

Distribution of Paris exhibitors across award categories. Table 3 shows how award categories relate to selection for Paris. Overall, 14.7% of firms were selected. Participation is strictly positive in all award categories and declines monotonically with award rank: 61.5% among gold medal recipients, 36.3% for silver, 16.0% for bronze, 7.0% for honorable mentions, and 6.8% for firms without an award. This monotonic pattern is consistent with our interpretation of the awards as ranking export potential, and with a selection committee that prioritized firms with higher export potential.

Table 3: Distribution of Firms Across 1896 Award Categories, and Participation in the Paris Exhibition

	Firms in sample		**Paris exhibitors**		**Non-Paris exhibitors**	
Award category	**Number**	**Share**	**Number**	**Share**	**Number**	**Share**
Gold medal	135	3.7%	83	61.5%	52	38.5%
Silver medal	405	11.2%	147	36.3%	258	63.7%
Bronze medal	1,007	27.9%	161	16.0%	846	84.0%
Honorable mention	1,122	31.0%	78	7.0%	1,044	93.0%
No award	945	26.1%	64	6.8%	881	93.2%
Total	3,614	100%	533	14.7%	3,081	85.3%

Industrial representativeness of Paris exhibitors. If Paris exhibitors were perfectly sorted by export potential, there would be little overlap between treated and untreated firms within award categories, undermining identification. It is therefore not sufficient that awards capture export potential: we also need deviations from perfect sorting that are unrelated to export potential. The government constraints described in Section 2, in particular the desire for industrial and regional balance, have the potential to generate such deviations.

Figure 2 provides evidence that industrial representation was a motive for deviating from perfect selection. Each bar plots, for a given industry, the difference between the industry's share among a particular group of firms and its share in the overall population of Budapest exhibitors. The sum of squared deviations (SSD) across industries summarizes how far each group's composition departs

[15] For details, see Appendix E.

from the population distribution.[16]

The blue bars show this deviation for gold medalists alone. The mismatch is substantial (SSD = 4,002): the food and beverage industry had a disproportionate share of high-potential firms, consistent with Hungary's early success in flour milling, while clothing was strongly under-represented, consistent with the concentration of textile and clothing production in the Austrian and Bohemian parts of the Habsburg empire, which left Hungarian manufacturers at a competitive disadvantage (Ránki, 1964).

Since only 135 firms received gold medals but the government sent 533 to Paris, positive selection alone would have implied sending both gold and silver medalists, who together numbered 540. The green bars show that expanding the pool to include silver medalists improves industrial balance, but substantial imbalances persist (SSD = 482).

The yellow bars show the actual composition of Paris participants. While perfect balance was not achieved, representativeness improved markedly: the SSD falls to 254, a 94% reduction relative to gold medalists and roughly half the gold-and-silver benchmark. This indicates that the selection committee actively balanced firms across industries, explaining why not only firms with the highest export potential were sent to Paris.

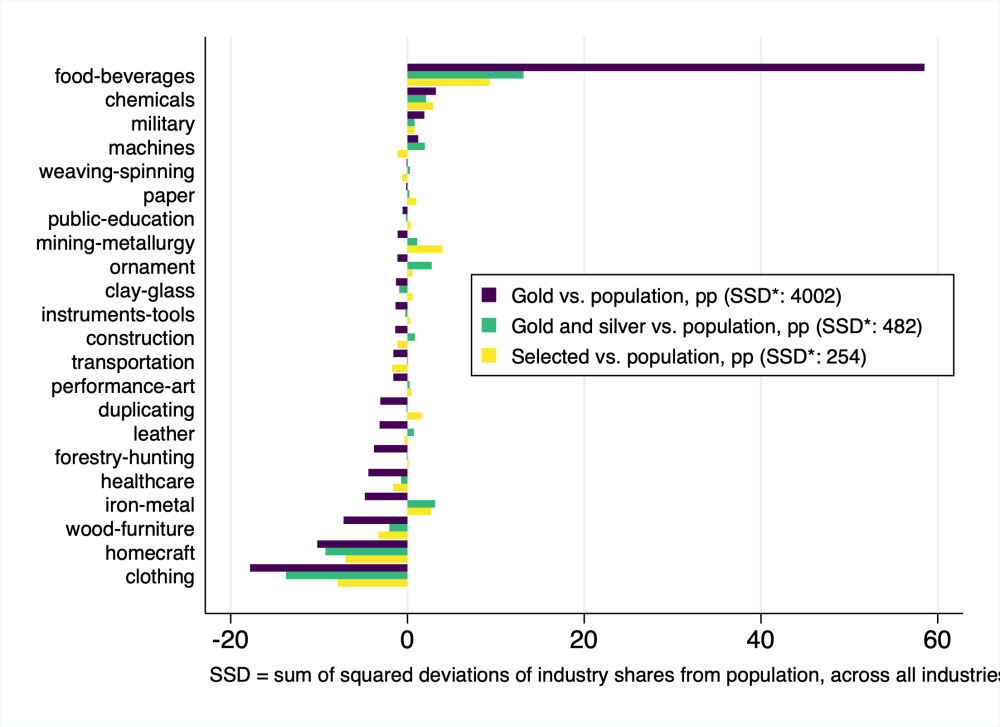


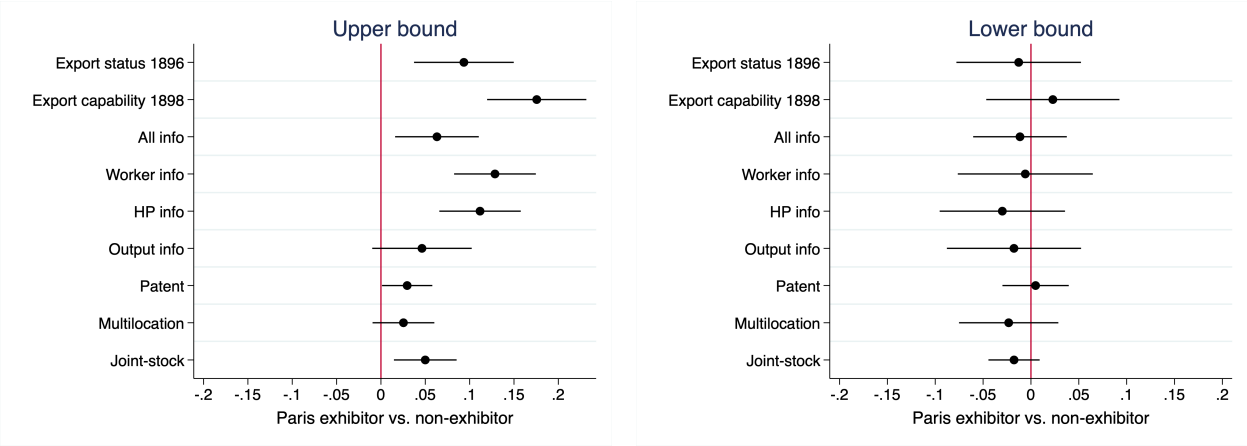
Figure 2: Deviation of industry shares from the population distribution: gold medalists, gold and silver medalists, and actual Paris participants

[16]The SSD is a descriptive measure of distributional distance, not a formal test statistic. We use it to summarize the degree of over- and under-representation across industries in a single number.

Balancing tests: pre-treatment characteristics. As discussed in Section 4.2, we conduct two sets of balancing checks by testing for differences in pre-treatment characteristics between treated and untreated firms within their respective comparison groups: within award classes for the upper bound (testing Assumption 1) and across adjacent award classes for the lower bound (testing Assumption 2). We implement these tests by estimating Equations (2) and (3) with pre-treatment characteristics as dependent variables. Results for indicator variables are reported in Figure 3, and for continuous variables in Figure 4.

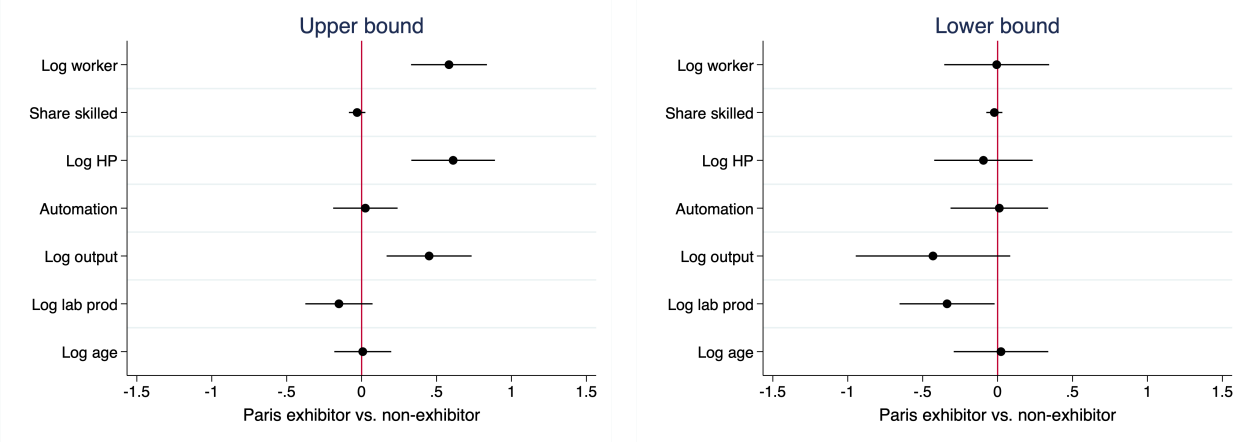
Figures 3a and 4a present the upper bound balancing tests. Each coefficient represents the mean difference in a pre-treatment characteristic between treated and untreated firms within the same award category. Consistent with Assumption 1, nearly all differences are positive, with several statistically significant. Figure 3a shows that within award categories, treated firms were more likely to export, to report export capability, to provide information on employment and machine capacity, and to be organized as joint-stock companies, all indicative of larger and more established firms. Among firms with detailed characteristics available, Figure 4a confirms that treated firms tend to be larger in employment, machine capacity, and output, while the share of skilled workers and firm age are balanced.

Figures 3b and 4b present the lower bound balancing tests, where each coefficient represents the mean difference between treated firms in a lower award category and untreated firms in the adjacent higher category. The results are strikingly different: there are no positive and significant differences, indicating that the positive selection evident in the upper bound comparison has disappeared. Point estimates are closer to zero or negative, supporting Assumption 2. If anything, the across-award comparison introduces negative selection, as intended by the lower bound construction.



(a) Upper bound (b) Lower bound

Figure 3: Balancing tests: pre-treatment indicator variables



(a) Upper bound (b) Lower bound

Figure 4: Balancing tests: pre-treatment continuous variables

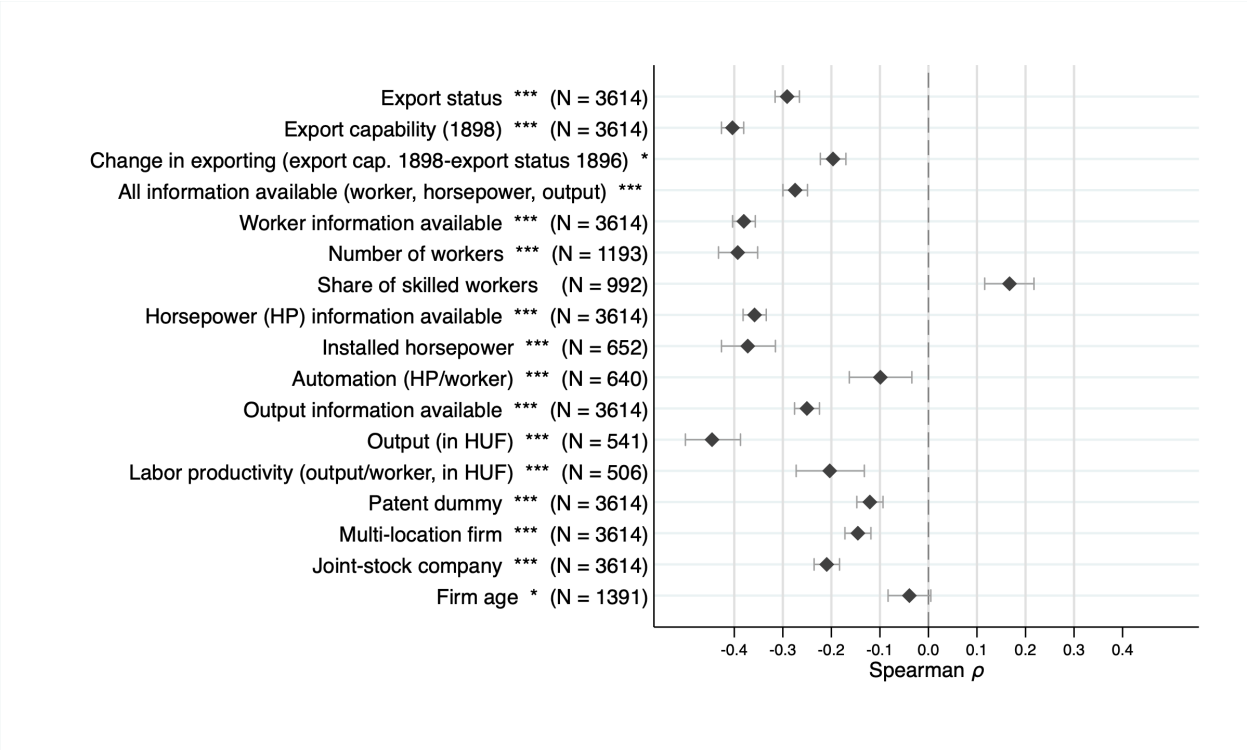
Validation test: Spearman rank correlation. As discussed in Section 4.2, Assumptions 1 and 2 jointly imply a complete ordering of unobserved export potential across subgroups defined by the interaction of award class and treatment status (Lemma 3). Figure 5 tests whether this ordering holds for observable pre-treatment characteristics using a Spearman rank correlation test. A negative correlation coefficient indicates that the characteristic declines with rank, consistent with the ordering predicted by the lemma.

The estimated coefficients are overwhelmingly consistent with the predicted ordering: all but one are negative, and nearly all are statistically significant. A positive or zero rank correlation can be rejected at the 1% level for almost all variables and at the 10% level for firm age. The sole exception is the share of skilled workers, which is positively correlated with rank. This is consistent with lower-ranked firms being less mechanized and therefore more reliant on craft skills, though the skill share is available for only 27% of observations.[17]

Overall, the rank correlation test supports the assumptions across a wide range of pre-determined firm characteristics.

[17]Indeed, Appendix Figure A.1 shows a negative correlation between the skill share and mechanization as well as firm size for the sample for which both measures are available.

Figure 5: Spearman rank correlation between pre-treatment firm characteristics and rank group



Notes: The figure plots Spearman rank correlation coefficients between pre-treatment firm characteristics and rank groups r as defined in Lemma 3. All variables are from 1896 unless indicated otherwise. The null hypothesis is $H_0 : \rho \geq 0$; the one-sided alternative is $H_1 : \rho < 0$. Stars denote one-sided p-values for this directional test: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Horizontal lines denote 90% confidence intervals constructed via Fisher z-transformation.

6 The Effect of the Paris Exhibition

Having presented evidence supporting Assumptions 1 and 2, we now estimate the effect of Paris Exhibition participation on the change in export capability and export status. We estimate two versions of the following regression, varying the fixed effects and sample as described in Section 4:

$$\Delta Y_i = \beta \, \mathrm{D}_i + \ FE_b + Y_{i,pre} + FE_p + \epsilon_i \tag{4}$$

where ΔY_i is the change in the outcome for firm i, D_i is an indicator for Paris Exhibition participation in 1900, and FE_b denotes bound-group fixed effects. For the upper bound estimation, FE_b consists of award category fixed effects, ensuring within-award comparisons on the full sample. For the lower bound estimation, FE_b consists of comparison group fixed effects that pair treated firms in lower award categories with untreated firms in higher categories, dropping treated firms in the top award category and untreated firms in the lowest. We control for the baseline level of the outcome $Y_{i,pre}$ and include product fixed effects FE_p to absorb differential export trends across product categories.

Export capability. Column (1) of Table 4 reports a simple regression of the change in the export capability indicator between 1898 and 1904 on the Paris indicator, controlling for initial export capability but without bound-group or product fixed effects. Conditional on initial export capability, Paris exhibitors are 19.1 percentage points more likely to report export capabilities in 1904.

Column (2) implements the upper bound estimation by adding award category and product fixed effects, reducing the estimate to 13.3 percentage points. Column (4) presents the lower bound, which falls further to 7.6 percentage points but remains statistically significant. Even the most conservative estimate thus provides evidence of a positive causal effect of trade fair participation. Relative to the baseline export capability share of 25% (Table 2), the lower bound implies an increase of roughly 30%, while the upper bound implies an increase of over 50%.

Column (3) provides an intermediate specification: it restricts the sample to firms used in the lower bound estimation while retaining the upper bound fixed effects. The coefficient is nearly identical to that in column (2), confirming that the reduction in the lower bound estimate is driven by the change in the control group rather than the change in sample composition.

Our estimation strategy also passes the second validation test specified in Lemma 4: the upper bound estimate (13.3pp) exceeds the lower bound estimate (7.6pp), as required by Assumptions 1 and 2. Failure of this inequality would cast doubt on at least one of the two assumptions. The fact that it holds, and that the gap between the bounds is plausible in magnitude, provides additional confidence in the identifying assumptions beyond the balancing tests and the Spearman rank correlation reported above.

Table 4: Treatment Effect on Change in Export Capability, 1904-1898

	(1)	(2)	(3)	(4)
		Upper bound		Lower bound
Paris exhibitor	0.191***	0.133***	0.136***	0.0760***
	(0.0266)	(0.0230)	(0.0271)	(0.0255)
Export capability 1898	Yes	Yes	Yes	Yes
Product fixed effects	No	Yes	Yes	Yes
Bound group dummies	No	Upper	Upper	Lower
Sample	Upper	Upper	Lower	Lower
Observations	3,614	3,614	2,650	2,650
R-squared	0.231	0.311	0.313	0.313

Notes: Sample of firms: columns (1) and (2) include all firms from the Budapest exhibition. Columns (3) and (4) only include firms used in the lower bound estimation sample. Standard errors are clustered by product. *** p<0.01, ** p<0.05, * p<0.1.

Export status. While it is encouraging that trade fair participation increased export capability, our ultimate interest lies in actual exporting. Table 5 examines the change in export status as the outcome, retaining the same column structure as Table 4.

The results confirm positive and significant effects on actual exporting: the upper bound estimate is 10.1 percentage points (column 2), while the lower bound is 4.2 percentage points (column 4), both statistically significant. As before, the naïve regression without fixed effects vastly overestimates the effect (column 1), and the reduction in the lower bound is driven by the change in control group rather than in sample composition. The estimates again pass the validation test of Lemma 4, with the upper bound exceeding the lower bound. Relative to the baseline export rate of 15% (Table 2), the lower bound implies an increase of roughly 28%, while the upper bound implies an increase of approximately 67%.

Table 5: Treatment Effect on Change in Export Status, 1905-1896

	(1)	(2) Upper bound	(3)	(4) Lower bound
Paris exhibitor	0.177***	0.101***	0.115***	0.0417*
	(0.0280)	(0.0220)	(0.0244)	(0.0240)
Export status 1896	Yes	Yes	Yes	Yes
Product fixed effects	No	Yes	Yes	Yes
Bound group dummies	No	Upper	Upper	Lower
Sample	Upper	Upper	Lower	Lower
Observations	3,614	3,614	2,650	2,650
R-squared	0.388	0.473	0.481	0.478

Notes: Sample of firms: columns (1) and (2) include all firms from the Budapest exhibition. Columns (3) and (4) only include firms used in the lower bound estimation sample. Standard errors are clustered by product. *** p<0.01, ** p<0.05, * p<0.1.

Placebo test: pre-treatment export performance. For exporting, our main outcome, we observe export-related performance twice before treatment: export status in 1896 and export capability in 1898. While these capture slightly different notions of exporting, they allow us to test whether treated and control firms exhibited differential pre-treatment dynamics. Table 6 estimates both specifications using the change between the two pre-period measures as the outcome, retaining the same column structure as Table 4. Because export status and export capability are not identical measures, we interpret this exercise as a placebo test rather than a formal pre-trend test.

The naïve estimate in column (1) shows that future Paris exhibitors were already experiencing faster export growth prior to the exhibition. Column (2) implements the upper bound specification and reveals a similar pattern, confirming that the upper bound estimates in Tables 4 and 5 overstate the causal effect.[18]

Importantly, the lower bound specification in column (4) shows no such differential: the coefficient falls to 1.8 percentage points and is statistically insignificant. Column (3) demonstrates that restricting the sample alone is insufficient to eliminate the pre-treatment difference in export growth; only the change in comparison group achieves this. The fact that the placebo effect vanishes in the lower bound specification provides strong evidence that comparing treated firms to untreated firms from higher award categories successfully eliminates positive selection bias, as our research design intends.

[18] A direct comparison of magnitudes to Tables 4 and 5 is not informative, as the pre-period outcome combines two different export measures.

Table 6: Placebo Test: Pre-treatment Export Performance, 1896–1898

	(1)	(2) Upper bound	(3)	(4) Lower bound
Paris exhibitor	0.197***	0.112***	0.107***	0.0180
	(0.0237)	(0.0238)	(0.0285)	(0.0340)
Export status 1896	Yes	Yes	Yes	Yes
Product fixed effects	No	Yes	Yes	Yes
Bound group dummies	No	Upper	Upper	Lower
Sample	Upper	Upper	Lower	Lower
Observations	3,614	3,614	2,650	2,650
R-squared	0.086	0.234	0.226	0.223

Notes: Sample of firms: columns (1) and (2) include all firms from the Budapest exhibition. Columns (3) and (4) only include firms used in the lower bound estimation sample. Standard errors are clustered by product. *** p<0.01, ** p<0.05, * p<0.1.

6.1 Comparison to Manski Bounds and Matching Estimators

As discussed in Section 4.3, our approach is related to Manski's partial identification framework (Manski, 2003). In Appendix C, we formally compare our identifying assumptions to the standard alternatives of bounded outcomes, MTR, MTS, and MIV. Here, we compare the bounds generated by each approach in our setting, illustrated in Figure 6.

We begin with the weakest assumption: that outcomes are bounded. For a binary outcome such as export status, $Y \in [0, 1]$. The Manski bounds set the unobserved counterfactual outcomes to their extreme feasible values: the lower bound assumes all treated units would have exported without treatment, while the upper bound assumes none would have. The first line in Figure 6 reports the resulting bounds: $[-0.435, 0.565]$ for export capability and $[-0.688, 0.311]$ for export status. With binary outcomes, these bounds always have width 1 and necessarily include zero, making them too wide for inference.

Researchers often impose two additional restrictions to tighten these bounds. *Monotone treatment selection (MTS)* assumes that treated units have weakly higher counterfactual outcomes than untreated units, while *monotone treatment response (MTR)* assumes that treatment weakly improves outcomes for each unit. The second line shows how MTR and MTS narrow the bounds: the lower bound becomes 0 by construction, and the upper bound reduces to the naïve comparison $\mathbb{E}[Y \mid D = 1] - \mathbb{E}[Y \mid D = 0]$, which ignores selection. The resulting intervals are [0, 0.390] for export capability and [0, 0.242] for export status. While substantially tighter, these bounds still include zero by construction and therefore cannot rule out null effects.

The third line applies an alternative refinement: using the award categories as a *Monotone Instrumental Variable* (MIV) in the sense of Manski and Pepper (2000). Despite exploiting the same covariate that our approach uses, the MIV bounds are wide, tightening the bounded-outcomes intervals only marginally: $[-0.435, 0.543]$ for export capability and $[-0.667, 0.291]$ for export status.

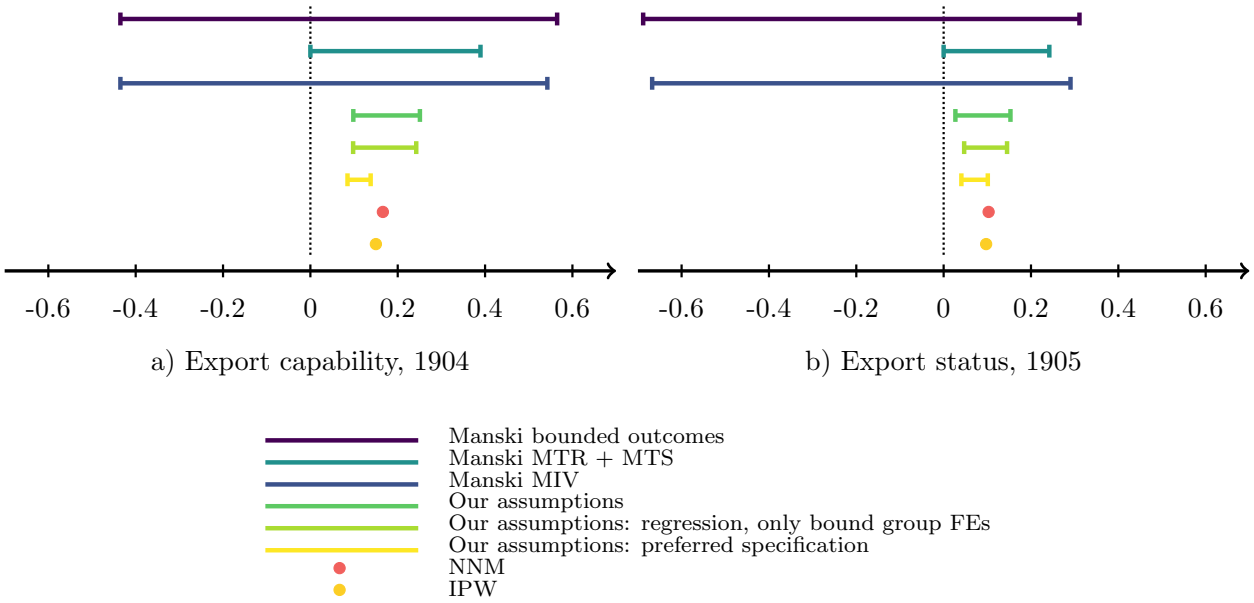


Figure 6: Comparison To Manski bounds and Matching Estimators

The fourth line applies our identifying assumptions (Assumptions 1 and 2) within the same partial-identification logic, computing subgroup-specific means weighted by treatment-group shares. Relative to MTS, our upper bound imposes positive selection *within* award classes rather than globally. Relative to MIV, our lower bound leverages outcomes for untreated units that we argue have weakly better no-treatment potential than the treated, restricting the comparison to specific subgroups rather than averaging across all firms in an award category. Because our assumptions neither nest within nor strictly strengthen MTR, MTS, or MIV, the resulting identified set need not be contained in these alternative bounds. Empirically, however, our bounds tighten substantially to [0.098, 0.251] for export capability and [0.027, 0.153] for export status, excluding zero in both cases.

The fifth line shows that implementing our assumptions via OLS with bound-group fixed effects, as described in Section 4.1, yields very similar results.[19] The sixth line reports our preferred specification, which adds controls for past exporting and product fixed effects as in Tables 4 and 5. These controls further narrow the bounds, primarily by lowering the upper bound.

The last two lines report point estimates from two conventional matching approaches that researchers commonly use to control for selection on observables. We assume that awards are unobserved and instead predict Paris participation using baseline firm characteristics alone (see Appendix Table B.3). Past export status, export capability, reporting horsepower information, and being a joint-stock company are significant predictors of participation. We then implement

[19]The regression-based implementation need not coincide exactly with the subgroup-mean construction because OLS aggregates across award classes using implicit weights proportional to within-award treatment variation.

nearest-neighbor matching (NNM, with 5 neighbors) and inverse probability weighting (IPW) based on these predicted propensities. The resulting point estimates lie at or above the upper bound of our preferred specification, suggesting that conventional matching on observables is insufficient to eliminate positive selection bias. This underscores the value of exploiting the award categories as a predetermined proxy for unobserved export potential.

6.2 Robustness: Alternative Upper Bounds

Assumption 1, which we use to derive upper bounds, formalizes *positive selection* within award classes: on average, the government selected firms with higher export potential within a given award category to attend Paris. This does not require that every treated firm has higher export potential than every untreated firm within the same award class; it is sufficient that positive selection occurs *on average* within each class. However, this assumption may be violated. Constraints on space, sectoral and regional representation, or political favoritism (e.g., lobbying) could attenuate positive selection or even reverse it. In the latter case, our "upper bound" could cease to be an upper bound.

To assess robustness, we therefore consider a weaker comparison that does not rely on within-award positive selection. Instead, we only require that treated firms in award class a have (weakly) higher export potential than untreated firms in the *next-lower* award class $a+1$. We formalize this as follows.

Assumption 3 (Weaker upper bound)**.** *Selected firms have (weakly) higher export potential than non-selected firms from the next-lower award class:*

$$\mathbb{E}[Y_i(0) \mid D_i = 1, A_i = a] \;\geq\; \mathbb{E}[Y_i(0) \mid D_i = 0, A_i = a+1], \quad \forall a \in \{1,2,3,4\}.$$

This assumption yields an alternative *upward-biased* comparison for each award class, and thus an alternative upper bound on the ATT.

Lemma 5 (Weaker upper bound)**.** *Define the alternative upper bound estimate for each award group a as*

$$\beta^{a,alt}_{upper} := \mathbb{E}[Y(1) \mid D = 1, A = a] - \mathbb{E}[Y(0) \mid D = 0, A = a+1],$$

and the average treatment effect on the treated for each award category a as

$$\tau^a := \mathbb{E}[Y(1) \mid D = 1, A = a] - \mathbb{E}[Y(0) \mid D = 1, A = a].$$

Under Assumption 3,

$$\beta^{a,alt}_{upper} \geq \tau^a,$$

i.e., $\beta^{a,alt}_{upper}$ is an upper bound on τ^a.

Proof.

$$\begin{aligned}
\beta_{\text{upper}}^{a,\text{alt}} &= E[Y(1) \mid D=1, A=a] - E[Y(0) \mid D=0, A=a+1] \\
&= E[Y(1) \mid D=1, A=a] - E[Y(0) \mid D=1, A=a] \\
&\quad + E[Y(0) \mid D=1, A=a] - E[Y(0) \mid D=0, A=a+1] \\
&= \tau^a + \underbrace{E[Y(0) \mid D=1, A=a] - E[Y(0) \mid D=0, A=a+1]}_{\geq 0}.
\end{aligned}$$

□

We implement this comparison by constructing "alternative upper-bound" comparison groups $h \in \{1, 2, 3, 4\}$, where group h contains treated firms in award class h and untreated firms in award class $h+1$. Comparing treated and untreated firms within the same comparison group yields an upward-biased estimate for each group h.

We aggregate these comparison-group-specific estimates using the regression

$$\Delta Y_i = \beta_{\text{upper}}^{\text{alt}} \cdot D_i \; + \; \sum_h \theta_h \cdot \mathbf{1}\{H_i = h\} \; + \; \epsilon_i, \tag{5}$$

where H_i denotes comparison-group membership. This specification changes the sample by excluding treated firms in the lowest award category ($A = 5$, since there is no next-lower control group) and untreated firms in the highest award category ($A = 1$, since there is no higher-ranked treated group).

Table 7: Treatment Effects with Alternative Upper Bounds under Weaker Assumptions

	Δ Export capability, 1904-1898		Δ Export status, 1905-1896	
	Alt. Upper	Lower	Alt. Upper	Lower
Paris exhibitor	0.206***	0.0760***	0.183***	0.0417*
	(0.0273)	(0.0255)	(0.0288)	(0.0240)
Initial outcome	Yes	Yes	Yes	Yes
Product fixed effects	Yes	Yes	Yes	Yes
Bound group dummies	Alt. Upper	Lower	Alt. Upper	Lower
Observations	3,498	2,650	3,498	2,650
R-squared	0.314	0.313	0.475	0.478

Notes: The samples refer to the alternative upper and lower bound sample, respectively. Standard errors are clustered by product. *** p<0.01, ** p<0.05, * p<0.1.

Table 7 reports the resulting alternative upper bounds alongside our baseline lower bounds. As before, we validate the procedure by verifying that the upper bounds exceed the lower bounds, which holds for all outcomes. In addition, because Assumption 3 is weaker than Assumption 1,

the alternative upper bounds should be weakly larger than our baseline upper bounds. This additional validation check is satisfied for our outcomes, providing further support for our award-group ordering.

6.3 Other Outcomes: Firm Survival and Employment Growth

While the main focus of this paper is the effect of trade fairs on exporting, in this section we report estimates for two additional outcomes available in our data: firm survival and employment growth.

Our baseline sample excludes 122 firms that exited between 1896 and 1904, as we estimate treatment effects conditional on survival. In Table 8, we examine the effect of Paris participation on exit directly by adding these firms back to the sample and using an indicator for firm exit as the outcome. The first two columns report the estimated bounds: participation in the Paris Exhibition significantly reduced exit probability.

Since no complete firm registry exists for this period, we cannot be certain that our exit records are comprehensive. As a complementary exercise, we re-estimate the export specifications on a sample of verified survivors using an additional source. Specifically, we located the 1906–1907 *Magyar Gyárak és Iparvallalatok Címtára* (Directory of Hungarian Factories and Industrial Enterprises), which lists a subset of firms, likely larger and more mechanized establishments. Firms appearing in this directory must have survived at least through 1906–1907. Using appearance in this directory as an outcome in columns 3 and 4, we find a positive and significant upper bound but a lower bound indistinguishable from zero. This pattern is consistent with the Paris Exhibition also improving survival prospects, though we cannot determine whether increased survival is a consequence of increased exporting or operates through other channels.[20]

[20] Appendix Table B.1 confirms positive upper and lower bounds on the treatment effect for this subset of surviving firms.

Table 8: Treatment Effects on Firm Exit and Survival

	Firm exit 1896-1904		Directory appearance 1906	
	Upper bound	Lower bound	Upper bound	Lower bound
Paris exhibitor	−0.0249***	−0.0371***	0.138***	0.0391
	(0.00819)	(0.0117)	(0.0303)	(0.0370)
Product fixed effects	Yes	Yes	Yes	Yes
Bound group dummies	Upper	Lower	Upper	Lower
Observations	3,736	2,742	3,614	2,650
R-squared	0.029	0.037	0.226	0.217

Notes: Bound group dummies refer to upper and lower bound sample restrictions. The sample for the first two columns contain all Budapest exhibitors in our analysis, the sample for the directory appearance does not contain firms with confirmed exit between 1896-1904. Standard errors are robust and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

For a subset of firms, we also observe employment after the exhibition (Table 2). Table 9 reports estimates using the change in log workers as the outcome. The upper bound on the treatment effect is 23.0% (column 2) and the lower bound is 16.3% (column 4), both statistically significant despite the smaller sample driven by the availability of employment data. An employment increase of 16–23% over eight years corresponds to an annual growth rate of roughly 1.9–2.6%. For a median firm with 35 workers in 1896, this translates into approximately 5–8 additional jobs by 1904.

Table 9: Treatment Effect on Employment Growth, 1896-1904

	(1)	(2)	(3)	(4)
		Upper bound		Lower bound
Paris exhibitor	0.227***	0.230***	0.162**	0.163**
	(0.0572)	(0.0637)	(0.0725)	(0.0776)
ln(workers), 1896	Yes	Yes	Yes	Yes
Product fixed effects	No	Yes	Yes	Yes
Bound group dummies	No	Upper	Upper	Lower
Sample	Upper	Upper	Lower	Lower
Observations	644	644	538	538
R-squared	0.104	0.188	0.177	0.175

Notes: Sample of firms: columns (1) and (2) include all firms from the Budapest exhibition where the number of workers is observable both in 1896 and 1904. Columns (3) and (4) include only firms used in the lower bound estimation sample. Standard errors are clustered by product. *** p<0.01, ** p<0.05, * p<0.1.

7 Competition and Market Thickness

Unlike interventions that provide individual firms with targeted access to new buyers or suppliers, trade fairs bring together large numbers of firms simultaneously. There are two potential forces at play: firms may find potential buyers but also compete for attention. To examine this, we estimate heterogeneous treatment effects based on the number of competitors firms faced at the trade fair as well as market thickness – the number of potential buyers.

To do this, for all firms in Paris, we leveraged the product subcategory[21] definition determined for all exhibitors. Some of the largest categories were "Wines and wine spirits" or "Mining, mining operations, and quarries", while examples for smaller ones include "Greenhouse Plants" or "Spinning and rope-making equipment and processes". The Paris catalog listed firms by these categories. There are 72 subcategories that included Hungarian firms and are relevant for private enterprises.

We construct two main measures of competition and buyers.

First, regarding competition, we measure the number of Hungarian competitors as the total number of Hungarian firms in the same Paris subcategory. Hungarian firms have known each other, and have little to gain in terms of business in Paris. A larger number of Hungarian exhibitors can intensify competition (more rivals in the same niche), but it could also have a positive effect if a stronger Hungarian presence helps attract more foreign buyers to that subcategory.

Second, we construct a buyers measure based on non-Hungarian exhibitors. Here we focus on the product categories where potential buyers of the firm's product are likely to be found. Using AI-assisted classification, we identify product categories that plausibly contain buyers for a given firm's product, explicitly excluding the firm's own product category. We then sum the number of exhibitors across all these "buyer" categories. This gives a measure of how many potential non-Hungarian buyers a firm is exposed to at the exhibition.

There is a large variation in both the number of Hungarian competitors (mean of 44, with a range of 3 to 157) and international buyers (mean of 3,149, with a range of 92 to 13,874). When ranking subcategories, there are few Hungarian firms in "Electrochemistry" or "Silks and silk fabrics", while domestic competition is intense in "Wines and wine spirits" or "Equipment and processes of food industries" In terms of international buyers, tough product categories are "Hunting weapons" or "Musical instruments", while we can expect many buyers in "Flour products and their derivatives", or "Mining, mining operations, and quarries."

We estimate the following regression, which is an extension of the regression described in equation (4) with the interaction terms for competition and potential buyers:

[21] Within each major industries of the time like Mines and metallurgy, there were several subcategories denoting finer product groups.

$$\Delta Y_i = \beta \, \mathrm{D}_i + \gamma \mathrm{D}_i \times \mathrm{Competition}_k + \delta \mathrm{D}_i \times \mathrm{Buyer}_k + FE_b + Y_{i,pre} + FE_p + \epsilon_i \quad (6)$$

where outcomes and fixed effects are defined as in equation (4). $\gamma \mathrm{D}_i \times \mathrm{Competition}_k$ represents the effect of competition in Paris subcategory k on the outcomes for firm i if i is a Paris exhibitor. $\delta \mathrm{D}_i \times \mathrm{Buyer}_k$ measures the effect of the number of potential buyers for Paris subcategory k for firm i if i is a Paris exhibitor.

Effects on export capability are presented in Table 10 using the only the lower bound estimation method. Competition and buyer market size are both measured as ln number of firms, and interacted with Paris indicator. The difference between the two columns is having all buyers in column (1) vs filtering out Hungarian buyers (in other product categories) in others.

Table 10: Effect of competition and potential buyers on change in export capability, 1898-1904

	(1)	(2)
Paris exhibitor * ln(Hungarian competition)	-0.0418*	-0.0428*
	(0.0249)	(0.0250)
Paris exhibitor * ln(number of buyers)	0.0414**	
	(0.0167)	
Paris exhibitor * ln(number of international buyers)		0.0434**
		(0.0175)
Product effects	Yes	Yes
Bound group dummies	Lower	Lower
Sample	Lower	Lower
Observations	2,721	2,721
R-squared	0.487	0.487

Notes: Only include firms used in the lower bound estimation sample. Standard errors are clustered by product. *** p<0.01, ** p<0.05, * p<0.1.

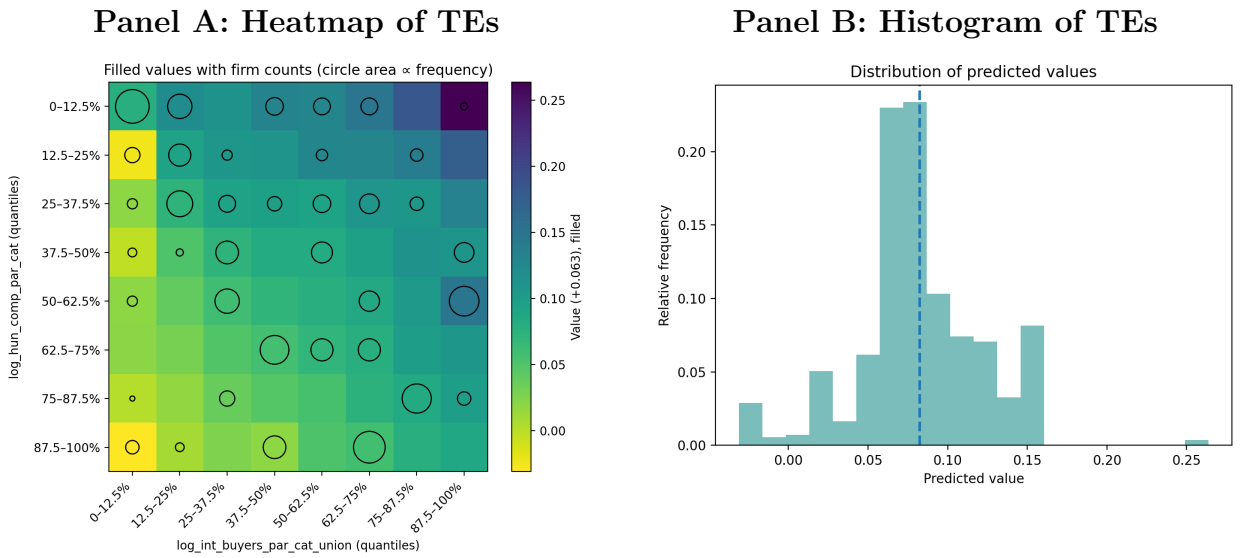
As Table 10 shows both mechanisms are at play: Hungarian firms that are part of product subcategories with many other Hungarian firms will benefit less of being in Paris, as attention is more divided and it is harder to make an impression. Conversely, Hungarian firms in product subcategories with many potential buyers benefit more.

Comparing firms at the bottom and top quartiles (the ln number of competitors at p25 is 2.77, and p75 is 4.11), the 1.34 log unit difference corresponds to a 5.74% (4.28*1.34) higher export

probability or two-thirds (66%) of the average treatment effect of 8.63% (Table 4. Repeating the same comparison for the ln number of international buyers (p25 is 7.34, and p75 is 8.21), the 0.87 log unit difference corresponds to a 3.7% (4.3*0.87) higher export probability or almost half (43%) of the average treatment effect. Both these variations are large, and for the upper bound estimates, they would be even larger.

Figure 7 shows the heterogeneity of the treatment effect. As both panels confirm, there is a large variation, the treatment varies between -0.02 and 0.25. There is only a very small segment of firms where the Paris exhibition has basically no effect (or even slightly negative point estimates), and a very small case for extremely high effect (where few Hungarian competitors and many potential buyers coincide). Still, there is a large variation between 5% and 15%, around the top right part of the joint distribution.

Figure 7: Heterogeneous treatment effects by competition and international buyers



Note. Using the distribution of firms by competition and international buyer (ln) values, we create an 8×8 grid based on a parametric model mapped onto the realized joint distribution, and estimate the heterogeneous treatment effect distribution by competition and international buyers. The left panel shows a heatmap of treatment effects; missing cells (with zero firms) are filled by iterative averaging of neighboring cells to obtain a smooth surface. Black hollow circles indicate the number of firms in each cell, with circle area proportional to cell frequency. The right panel shows a histogram of treatment-effect values, with the weighted mean indicated by a dashed line.

These results suggest that being in Paris has a heterogeneous treatment effect across product subcategories: firms with fewer domestic peers and more potential buyers are the ones able to benefit more in terms of export opportunities. We also learned that a small segment of firms did not benefit at all. This result is important as it indicates search and matching frictions are important and competition and large buyer base could alter the cost-benefit analysis of being at a large event.

These patterns of heterogeneity provide insight into the specific nature of search and matching

frictions that trade fairs address. Trade fairs reduce discovery and inspection costs by concentrating buyers and sellers in a single location for a limited time, effectively lowering the search radius. However, our results reveal two important dimensions of how this mechanism operates in practice. First, the negative effect of domestic competition demonstrates a congestion externality: when multiple similar sellers compete for buyers' limited attention and time, the individual benefit of reduced search costs diminishes. This suggests that search frictions are not eliminated. Second, the positive effect of potential buyer concentration indicates that matching frictions remain even after search costs are reduced. A thicker buyer market increases both the probability of finding a suitable match and the expected quality of that match (Lenoir et al., 2023).

8 Conclusion

This paper provides causal evidence that trade fairs can effectively promote exports by reducing search and matching frictions. Using the 1900 Paris Exhibition as a natural experiment, we show that participation increased Hungarian firms' export probability by 4-10 percentage points and generated 16-23% employment growth over the following decade. Export probabilities were highest for firms in thick markets (many potential buyers) but with few domestic competitors.

Our methodological contribution, using predetermined quality measures and selection constraints to bound treatment effects, addresses a fundamental challenge in evaluating targeted policies. When perfect randomization is infeasible, our approach shows how researchers can leverage institutional details to obtain credible causal estimates. In comparison to existing causal bands, our method offer a substantially narrower range.

Our findings carry implications for contemporary export promotion policy. Despite substantial public investment in trade fair subsidies, rigorous causal evidence on their effectiveness remains limited. We show that participation can generate substantial and persistent gains in both export probability and employment. However, the heterogeneity we document is equally important for policy design: moving from the 25th to the 75th percentile of domestic competitors cuts the treatment effect by two-thirds, while the same move along the buyer margin increases it by nearly half. This means that simply maximizing participation may be counterproductive; sending too many similar firms dilutes the benefit for each. Instead, export promotion agencies should limit the number of supported firms per product category and prioritize sectors where a large international buyer base can be reached. More broadly, our findings favor well-organized large events that concentrate many potential buyers with relatively few supported firms per niche.

References

Alfaro, Laura, Cathy Ge Bao, Maggie X. Chen, Junjie Hong, and Claudia Steinwender, "Omnia Juncta in Uno: Foreign Powers and Trademark Protection in Shanghai's Concession Era," NBER Working Paper 29721 2022.

Allen, Treb, "Information frictions in trade," *Econometrica*, 2014, *82* (6), 2041–2083.

Arkolakis, Costas, Federico Huneeus, and Yuhei Miyauchi, "Production Network Formation, Trade, and Welfare," 2025.

Atkin, David, Amit K. Khandelwal, and Adam Osman, "Exporting and Firm Performance: Evidence from a Randomized Experiment," *The Quarterly Journal of Economics*, 2017, *132* (2), 551–615.

Bergquist, Lauren F, Craig McIntosh, and Meredith Startz, "Search costs, intermediation, and trade: Experimental evidence from Ugandan agricultural markets," 2024.

Bernard, Andrew B., Andreas Moxnes, and Karen Helene Ulltveit-Moe, "Two-Sided Heterogeneity and Trade," *The Review of Economics and Statistics*, 07 2018, *100* (3), 424–439.

Bernard, Andrew B, Andreas Moxnes, and Yukiko U Saito, "Production networks, geography, and firm performance," *Journal of Political Economy*, 2019, *127* (2), 639–688.

Biesebroeck, Johannes Van, Emily Yu, and Shenjie Chen, "The impact of trade promotion services on Canadian exporter performance," *Canadian Journal of Economics/Revue canadienne d'économique*, 2015, *48* (4), 1481–1512.

Blundell, Richard, Amanda Gosling, Hidehiko Ichimura, and Costas Meghir, "Changes in the Distribution of Male and Female Wages Accounting for Employment Composition Using Bounds," *Econometrica*, 2007, *75* (2), 323–363.

Broocks, Annette and Johannes Van Biesebroeck, "The impact of export promotion on export market entry," *Journal of International Economics*, 2017, *107*, 19–33.

Buus, Magnus T, Jakob R Munch, Joel Rodrigue, and Georg Schaur, "Do Export Support Programs Affect Prices, Quality, Markups, and Marginal Costs? Evidence From a Natural Policy Experiment," *Review of Economics and Statistics*, 2025, *107* (1), 172–187.

Cai, Jing, Wei Lin, and Adam Szeidl, "Firm-to-Firm Referrals," 2024.

Carballo, Jerónimo and Christian Volpe Martincus, "Is export promotion effective in developing countries? Firm-level evidence on the intensive and the extensive margins of exports," *Journal of International Economics*, 2008, *76* (1), 89–106.

Demir, Banu, Ana Cecilia Fieler, Daniel Yi Xu, and Kelly Kaili Yang, "O-ring production networks," *Journal of Political Economy*, 2024, *132* (1), 200–247.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Xianglong Kong, Magne Mogstad, and Felix Tintelnot, "Endogenous production networks with fixed costs," *Journal of International Economics*, 2023, *145*, 103841.

Eaton, Jonathan, David Jinkins, James R Tybout, and Daniel Xu, "Two-sided search in international markets," Technical Report, National Bureau of Economic Research 2022.

_ , Samuel S Kortum, and Francis Kramarz, "Firm-to-Firm Trade: Imports, Exports, and the Labor Market," Working Paper 29685, National Bureau of Economic Research January 2022.

Görg, Holger, Michael Henry, and Eric Strobl, "Grant support and exporting activity," *The review of economics and statistics*, 2008, *90* (1), 168–174.

Hanlon, W Walker, "The persistent effect of temporary input cost advantages in shipbuilding,

1850 to 1911," *Journal of the European Economic Association*, 2020, *18* (6), 3173–3209.

Hiller, Sanne, "Does export promotion work in Denmark? Evidence from a matching approach," 2012.

Ho, Kate and Adam M Rosen, "Partial identification in applied research: benefits and challenges," *Advances in Economics and Econometrics*, 2017, *2*, 307–359.

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

_ and Nolan H Miller, "Market integration, demand, and the growth of firms: Evidence from a natural experiment in India," *American Economic Review*, 2018, *108* (12), 3583–3625.

Juhász, Réka and Claudia Steinwender, "Spinning the web: The impact of ICT on trade in intermediates and technology diffusion," Technical Report, National Bureau of Economic Research 2018.

_ and _ , "Industrial Policy and the Great Divergence," *Annual Review of Economics*, 2024, *16.*

Juhász, Réka, "Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade," *American Economic Review, 108 (11): 3339-76.*, 2018.

Klein, Alexander, Max-Stephan Schulze, and Tamás Vonyó, "How Peripheral was the Periphery? Industrialization in East Central Europe since 1870," in "The Spread of Modern Industry to the Periphery since 1871," Oxford University Press, 02 2017.

Krolikowski, Pawel M and Andrew H McCallum, "Goods-market frictions and international trade," *Journal of International Economics*, 2021, *129*, 103411.

Kun, Gergő, "Industrial Development Policies in the Austro-Hungarian Empire, Focusing on Hungary's Industry from the Austro-Hungarian Compromise of 1867 to the World War I," *SEA–Practical Application of Science*, 2019, *7* (19), 55–64.

Lederman, Daniel, Marcelo Olarreaga, and Lucy Payton, "Export Promotion Agencies: Do They Work?," *Journal of Development Economics*, 2010, *91* (2), 257–265.

Lee, David S, "Training, wages, and sample selection: Estimating sharp bounds on treatment effects," *The Review of Economic Studies*, 2009, *76* (3), 1071–1102.

Lenoir, Clémence, Julien Martin, and Isabelle Mejean, "Search frictions in international goods markets," *Journal of the European Economic Association*, 2023, *21* (1), 326–366.

Lim, Kevin, "Firm-to-firm Trade in Sticky Production Networks," 2017 Meeting Papers 280, Society for Economic Dynamics 2017.

Manski, Charles F, *Partial identification of probability distributions*, Springer, 2003.

Manski, Charles F. and John V. Pepper, "Monotone Instrumental Variables: With an Application to the Returns to Schooling," *Econometrica*, 2000, *68* (4), 997–1010.

Martincus, Christian Volpe and Jerónimo Carballo, "Beyond the average effects: The distributional impacts of export promotion programs in developing countries," *Journal of Development Economics*, 2010, *92* (2), 201–214.

McKinsey & Company, "With Artificial Intelligence, Find New Suppliers in Days, Not Months," 2021. Authors: Mark Baptista, Mike Gordon, Christine Herrmann, and Andrew Pratt. Accessed: 2025-08-01.

Miyauchi, Yuhei, "Matching and Agglomeration: Theory and Evidence from Japanese Firm-to-Firm Trade," *Econometrica*, 2024, *92* (6), 1869–1905.

Moser, Petra, "How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World's Fairs," *American Economic Review*, September 2005, *95* (4), 1214–1236.

_ , "Innovation without Patents: Evidence from World's Fairs," *The journal of law economics*, 2 2012, *55* (1), 43–74.

Munch, Jakob and Georg Schaur, "The effect of export promotion on firm-level performance," *American Economic Journal: Economic Policy*, 2018, *10* (1), 357–387.

Munch, J.R., M. Buus, J. Rodrigue, and G. Schaur, "Do Export Support Programs affect Prices, Quality, Markups and Marginal Costs? Evidence from a Natural Policy Experiment," *Review of Economics and Statistics*, forthcoming.

Ránki, György, "Problems of the development of Hungarian industry, 1900–1944," *The Journal of Economic History*, 1964, *24* (2), 204–228.

Romero, Dario A, "An Empire Lost: Spanish Industry and the Effect of Colonial Markets on Trade and Innovation," 2023.

Schminke, Annette and Johannes Van Biesebroeck, "Using export market performance to evaluate regional preferential policies in China," *Review of World Economics*, 2013, *149* (2), 343–367.

Schulze, Max-Stephan, "Patterns of growth and stagnation in the late nineteenth century Habsburg economy," *European Review of Economic History*, 2000, *4* (3), 311–340.

_ , "Origins of catch-up failure: comparative productivity growth in the Habsburg Empire, 1870–1910," *European Review of Economic History*, 2007, *11* (2), 189–218.

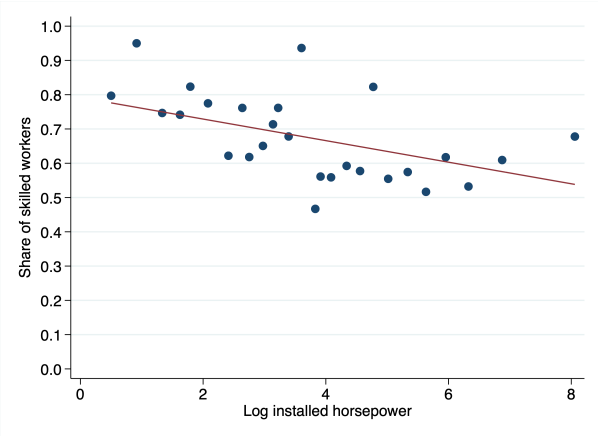
Steinwender, Claudia, "Real Effects of Information Frictions: "When the States and the Kingdom became United"," *American Economic Review*, 2018, *108* (3), 657–696.

Székely, Miklós, "The Resetting of the Main Historical Group from the Millennium Exhibition to the Paris Universal Exhibition of 1900," *Ephemeral Architecture in Central Eastern Europe in the 19th and 20th centuries*, 2015, pp. 33–50.

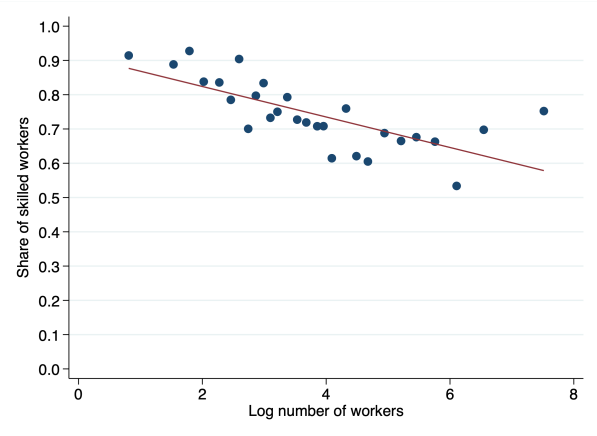
Tamer, Elie, "Partial identification in econometrics," *Annual Review of Economics*, 2010, *2* (1), 167–195.

Tamilia, Robert, "World's Fairs and the Department Store 1800s to 1930s," in "Proceedings of the Conference on Historical Analysis and Research in Marketing," Vol. 13 2007, pp. 228–240.

A Appendix Figures



(a) Skill share and mechanization



(b) Skill share and employment

Figure A.1: Relationship between skill share and mechanization or employment

Notes: Each dot represents the mean share of skilled workers within a quantile bin of log horsepower (Figure A.1a) and log number of workers (Figure A.1b). The line is a linear fit through the bin means.

B Appendix Tables

Table B.1: Treatment Effects for Verified Survivors

	Δ Export capability, 1904–1898		Δ Export status, 1905–1896	
	Upper bound	Lower bound	Upper bound	Lower bound
Paris exhibitor	0.109***	0.104**	0.170***	0.142***
	(0.0312)	(0.0416)	(0.0412)	(0.0506)
Export status 1898	Yes	Yes		
Export status 1896			Yes	Yes
Product fixed effects	Yes	Yes	Yes	Yes
Bound group dummies	Upper	Lower	Upper	Lower
Observations	736	614	736	614
R-squared	0.426	0.432	0.511	0.523

Notes: All columns include only firms from the Budapest exhibition that also appear in the 1906–1907 Directory of Hungarian Factories and Industrial Enterprises. Standard errors are clustered by product. *** p<0.01, ** p<0.05, * p<0.1.

Table B.2: Summary statistics by treatment and control group

		Share	Mean			
Variable	N	TG	TG	CG	Diff	t-stat
Baseline characteristics (1896, unless indicated otherwise)						
Export status	3,614	0.15	0.360	0.117	-0.243	-11.23***
Export capability (1898)	3,614	0.15	0.572	0.199	-0.373	-16.50***
Change in exporting (export cap. 1898-export status 1896)	3,614	0.15	0.212	0.081	-0.131	-6.01***
All information available (worker, horsepower, output)	3,614	0.15	0.300	0.102	-0.199	-9.64***
Worker information available	3,614	0.15	0.625	0.279	-0.346	-15.37***
Number of workers	1,193	0.28	364.650	97.962	-266.688	-3.87***
Share of skilled workers	992	0.28	0.700	0.779	0.078	3.19***
Horsepower (HP) information available	3,614	0.15	0.447	0.134	-0.312	-13.93***
Installed horsepower	652	0.37	393.649	127.876	-265.774	-2.79***
Automation (HP/worker)	640	0.37	0.969	0.884	-0.084	-0.68
Output information available	3,614	0.15	0.317	0.121	-0.196	-9.34***
Output (in HUF)	541	0.31	1.43 mil	350,000	-1.08 mil	-5.13***
Labor productivity (output/worker, in HUF)	506	0.33	6,073	6,935	861.828	0.37
Patent dummy	3,614	0.15	0.098	0.039	-0.059	-4.42***
Multi-location firm	3,614	0.15	0.129	0.055	-0.075	-4.96***
Joint-stock company	3,614	0.15	0.141	0.021	-0.119	-7.80***
Firm age	1,391	0.24	27.355	25.519	-1.836	-1.04
Post-treatment (1904, unless indicated otherwise)						
Export status (1905)	3,614	0.15	0.321	0.070	-0.251	-12.08***
Export capability	3,614	0.15	0.585	0.183	-0.402	-17.89***
Number of workers	796	0.36	372.180	89.312	-282.868	-4.04***

Table B.3: First Stage: Nearest Neighbor Matching and Propensity Score Estimation

	(1) Paris exhibitor
Export status indicator (1896)	0.281*
	(0.152)
Export capability indicator (1898)	1.078***
	(0.169)
Horsepower info available	0.604***
	(0.167)
Worker information available	0.160
	(0.177)
Patent indicator	0.230
	(0.210)
Joint-stock company	0.963***
	(0.212)
Output information available	−0.233
	(0.159)
Multi-location firm	0.192
	(0.186)
Observations	3,565
Product FE	Yes

Notes: Logit estimation. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table B.4: Number of Exhibitors – Millennial Exhibition, Budapest 1896

Group No.	Name	No. of Exhibitors	Floor Area (m^2) Official	Private	Total	Outdoor Area (m^2)
a) Historical Exhibition		880	7,500	—	7,500	300
b) Contemporary Exhibition						
I/a	Fine Arts	269	3,600	—	3,600	—
I/b	Performing Arts	65	980	—	980	—
II.	Public Education & Culture	105	600	1,880	2,480	—
III.	Education	697	2,580	160	2,740	—
IV.	Public Health & Child Welfare	287	1,780	810	2,590	—
V.	Trade, Finance & Credit[1]	344	—	1,780	1,780	—
VI.	Agriculture	3,175	8,390	1,310	9,700	2,600
VII.	Forestry	214	1,500	90	1,590	7,500
VIII/a	Mining & Metallurgy	72	3,350	1,750	5,100	50
VIII/b	Iron & Metal Industry	231	1,760	90	1,850	300
IX.	Machinery & Engineering	269	13,380	4,050	17,430	500
X.	Transport, Navigation & Maritime	175	5,600	730	6,330	2,400
XI.	Construction Industry	286	2,200	1,340	3,540	500
XII.	Wood, Furniture & Decorative Arts	285	4,400	330	4,730	—
XIII.	Ceramics & Glass Industry	108	1,140	—	1,140	75
XIV.	Leather, Spinning, Weaving & Clothing	997	5,970	—	5,970	—
XV.	Paper & Printing Industry	191	—	2,250	2,250	—
XVI.	Gold, Silver & Decorative Goods	83	1,150	—	1,150	—
XVII.	Military Affairs[2]	44	3,110	450	3,560	2,500
XVIII.	Chemical Industry	114	950	860	1,810	—
XIX.	Food Products as Industrial Goods	323	990	3,180	4,170	—
XX.	Cottage Industry, Ethnography & Amateur Exh.	436	4,700	—	4,700	—
Total		**9,650**	**75,630**	**21,060**	**96,690**	**16,725**

[1] Organized ex officio.
[2] Joint Army, Navy, Home Guard, Gendarmerie, Red Cross. Organized ex officio.
[3] The floor area of village houses is included in this figure.

C Manski bounds

This appendix reviews Manski's partial identification approach and shows how our award-based bounds tighten the bounds.

C.1 Definition of Manski bounds

Our object of interest is the average treatment effect on the treated (ATT),

$$\text{ATT} = \mathbb{E}[Y(1) \mid D = 1] - \mathbb{E}[Y(0) \mid D = 1], \tag{C.1}$$

where we observe only

$$Y = DY(1) + (1 - D)Y(0), \quad D \in \{0, 1\},$$

and we do not impose random assignment, selection-on-observables, instrumental variables, or other identifying assumptions. As a result, the counterfactual mean outcome for treated units,

$$\mathbb{E}[Y(0) \mid D = 1],$$

is not identified from the data. In our setting, Y is a dummy outcome (e.g., export status), so $Y \in [0, 1]$.

The key insight of Manski (2003) is that even without additional identifying assumptions, one can bound treatment effects using only the fact that outcomes are bounded.[22] For a binary outcome, the natural bounds are 0 and 1.

Assumption C.1 (Bounded outcomes)**.** *The potential outcomes are bounded in* $[0, 1]$*:*

$$0 \leq Y_i(0) \leq 1 \quad \textit{and} \quad 0 \leq Y_i(1) \leq 1 \qquad \textit{for all } i.$$

The Manski bounds for the ATT follow immediately by combining assumption C.1 with equation (C.1). The lower bound is obtained by setting the counterfactual to its largest feasible value (1), while the upper bound sets it to its smallest feasible value (0).

Lemma C.1 (Manski bounds)**.** *Define the upper and lower bounds on the ATT as*

$$\overline{ATT}_{Manski} = \mathbb{E}[Y(1) \mid D = 1] - 0,$$
$$\underline{ATT}_{Manski} = \mathbb{E}[Y(1) \mid D = 1] - 1.$$

Under assumption C.1,

$$\underline{ATT}_{Manski} \leq ATT \leq \overline{ATT}_{Manski}.$$

Proof.

$$\mathbb{E}[Y(1) \mid D = 1] - 1 \leq \mathbb{E}[Y(1) \mid D = 1] - \mathbb{E}[Y(0) \mid D = 1] \leq \mathbb{E}[Y(1) \mid D = 1] - 0.$$

□

That is, any value of the ATT outside $[\underline{\text{ATT}}_{\text{Manski}}, \overline{\text{ATT}}_{\text{Manski}}]$ is incompatible with the data. These bounds are sharp. Because these assumptions are extremely weak, the resulting bounds can

[22] Manski (2003) typically presents the argument for the ATE. Since our focus is the ATT, we state the bounds for $\mathbb{E}[Y(0) \mid D = 1]$ and hence for the ATT.

be wide. With a binary outcome, the width equals 1, so the ATT bounds necessarily include zero. This is often too coarse for inference, for instance when the goal is to rule out null effects.

C.2 Tightening the Manski bounds: Monotone treatment response (MTR) and monotone treatment selection (MTS)

Two common additional restrictions are the following. First, *monotone treatment selection (MTS)* assumes that treated units have weakly higher counterfactual outcomes than untreated units.

Assumption C.2 (Monotone Treatment Selection (MTS))**.**

$$\mathbb{E}[Y(0) \mid D = 0] \leq \mathbb{E}[Y(0) \mid D = 1].$$

Second, *monotone treatment response (MTR)* assumes that treatment weakly improves outcomes for each treated unit.

Assumption C.3 (Monotone Treatment Response (MTR))**.**

$$\mathbb{E}[Y(0) \mid D = 1] \leq \mathbb{E}[Y(1) \mid D = 1].$$

Together, these restrictions tighten the identified set for the ATT:

Lemma C.2 (Manski bounds with MTS and MTR)**.** *Define the upper bound using MTS and the lower bound using MTR for the counterfactuals:*

$$\overline{ATT}_{MTS} = \mathbb{E}[Y(1) \mid D = 1] - \mathbb{E}[Y(0) \mid D = 0],$$
$$\underline{ATT}_{MTR} = \mathbb{E}[Y(1) \mid D = 1] - \mathbb{E}[Y(1) \mid D = 1] = 0.$$

Under assumptions C.2 and C.3,

$$\underline{ATT}_{MTR} \leq ATT \leq \overline{ATT}_{MTS}.$$

The upper bound under MTS, $\overline{\text{ATT}}_{\text{MTS}}$, is often referred to as the *naive* comparison, since it effectively compares treated outcomes to untreated outcomes without accounting for selection.

C.3 Comparison to our bounds

Our approach extends the Manski approach by invoking assumptions 1 and 2 to derive the bounds. To see this, rewrite the ATT by conditioning on award class A:

$$\begin{aligned}\text{ATT} &= \mathbb{E}[Y(1) \mid D = 1] - \mathbb{E}[Y(0) \mid D = 1] \\ &= \sum_{a \in \mathcal{A}} \Big(\mathbb{E}[Y(1) \mid D = 1, A = a] - \mathbb{E}[Y(0) \mid D = 1, A = a] \Big) \Pr(A = a \mid D = 1), \end{aligned} \tag{C.2}$$

where $\Pr(A = a \mid D = 1)$ is the share of treated units in award category a.

Lemma C.3 (Our bounds)**.** *Define the upper bound using assumption 1 for the counterfactual,*

$$\overline{ATT} = \sum_{a \in \mathcal{A}} \Big(\mathbb{E}[Y(1) \mid D = 1, A = a] - \mathbb{E}[Y(0) \mid D = 0, A = a] \Big) \Pr(A = a \mid D = 1),$$

and the lower bound using assumption 2 for the counterfactual,

$$\underline{ATT} = \sum_{a \in \mathcal{A} \setminus \{1\}} \Big(\mathbb{E}[Y(1) \mid D = 1, A = a] - \mathbb{E}[Y(0) \mid D = 0, A = a - 1] \Big) \widetilde{\Pr}(A = a \mid D = 1)$$

where $\widetilde{\Pr}(A = a \mid D = 1) \equiv \Pr(A = a \mid D = 1, A \neq 1) = \frac{\Pr(A=a|D=1)}{\Pr(A \neq 1|D=1)}$*, i.e., it rescales weights to account for the fact that there is no counterfactual for the top award class. Assuming that the treatment effect for the excluded top award class is at least as large as the lower bound on the common support, and Assumptions 1 and 2,*

$$\underline{ATT} \leq ATT \leq \overline{ATT}.$$

Our Assumption 2 on the lower bound differs from both Manski and MTR in that we use the counterfactual for a group of untreated units that we assume to have better outcomes without treatment than the treated. Because the lower-bound construction relies on comparisons to the next-lower award category, it cannot be applied to the top award class. We therefore report bounds for the ATT among treated units with common support, as in trimming-based matching approaches. To interpret these bounds as applying to the overall ATT, we additionally assume that the treatment effect for the excluded top award class is at least as large as the common-support lower bound.

Our Assumption 1 on the upper bound differs from MTS in that it imposes Assumption C.2 only on within-award group comparisons.

C.4 Manski MIV bounds

Another refinement in the literature is to use an additional covariate that is believed to be related to potential outcomes as a Manski *Monotone Instrumental Variable (MIV)* (Manski and Pepper, 2000). A natural application in our setting is to use award categories A_i as the Manski MIV. Recall that $A_i \in \mathcal{A} = \{1, 2, 3, 4, 5\}$, where $A_i = 1$ indicates the highest export potential.[23] It is natural to assume that firms in better award categories have (weakly) higher export potential:

Assumption C.4 (Manski MIV)**.** *If* $a_1 \leq a_2$*, then*

$$\mathbb{E}[Y(0) \mid A = a_1] \geq \mathbb{E}[Y(0) \mid A = a_2] \quad \textit{and} \quad \mathbb{E}[Y(1) \mid A = a_1] \geq \mathbb{E}[Y(1) \mid A = a_2].$$

The following lemma illustrates how this assumption can be used to construct bounds for the counterfactual outcomes, which can then be used to construct bounds for the treatment effect on the treated ATT:[24]

Lemma C.4 (Manski MIV bounds)**.** *Assume that outcomes are bounded as in assumption C.1, and assume* $\Pr(D = 1 \mid A = a) > 0$ *for all* $a \in \mathcal{A}$ *with* $\Pr(A = a \mid D = 1) > 0$*. Define the*

[23] Notice that the ordering of covariates differs from the classical approach in Manski and Pepper (2000): in our setting, a smaller a corresponds to higher potential outcomes, for example, Gold has a smaller a than Silver.

[24] Notice that since we are only considering ATT bounds, only the restriction on $\mathbb{E}[Y(0) \mid A = a_1]$ in Assumption C.4 will be used.

award-specific bounds on the counterfactual outcome for the treated as

$$\overline{\mathbb{E}}[Y(0) \mid D=1, A=a] \equiv \min\Bigg\{1,\ \frac{1}{\Pr(D=1 \mid A=a)} \cdot \Bigg[\inf_{\tilde{a}\leq a}\Big(\Pr(D=1 \mid A=\tilde{a}) + \mathbb{E}[Y(0) \mid D=0, A=\tilde{a}]\Pr(D=0 \mid A=\tilde{a})\Big) - \mathbb{E}[Y(0) \mid D=0, A=a]\Pr(D=0 \mid A=a)\Bigg]\Bigg\},$$

$$\underline{\mathbb{E}}[Y(0) \mid D=1, A=a] \equiv \max\Bigg\{0,\ \frac{1}{\Pr(D=1 \mid A=a)} \cdot \Bigg[\sup_{\tilde{a}\geq a}\Big(\mathbb{E}[Y(0) \mid D=0, A=\tilde{a}]\Pr(D=0 \mid A=\tilde{a})\Big) - \mathbb{E}[Y(0) \mid D=0, A=a]\Pr(D=0 \mid A=a)\Bigg]\Bigg\}.$$

Define

$$\overline{ATT}_{IV} \equiv \sum_{a\in\mathcal{A}}\Big(\mathbb{E}[Y(1) \mid D=1, A=a] - \underline{\mathbb{E}}[Y(0) \mid D=1, A=a]\Big)\Pr(A=a \mid D=1),$$
$$\underline{ATT}_{IV} \equiv \sum_{a\in\mathcal{A}}\Big(\mathbb{E}[Y(1) \mid D=1, A=a] - \overline{\mathbb{E}}[Y(0) \mid D=1, A=a]\Big)\Pr(A=a \mid D=1).$$

Under assumption C.4,

$$\underline{ATT}_{IV} \leq ATT \leq \overline{ATT}_{IV}.$$

Proof. By the law of total expectation,

$$\mathbb{E}[Y(0) \mid A=a] = \mathbb{E}[Y(0) \mid D=1, A=a]\Pr(D=1 \mid A=a) + \mathbb{E}[Y(0) \mid D=0, A=a]\Pr(D=0 \mid A=a). \tag{C.3}$$

Upper bound on the counterfactual. In this expression, only $\mathbb{E}[Y(0) \mid D=1, A=a]$ is unobserved, but since $Y(0) \leq 1$, we have $\mathbb{E}[Y(0) \mid D=1, A=a] \leq 1$. Therefore:

$$\mathbb{E}[Y(0) \mid A=a] \leq \Pr(D=1 \mid A=a) + \mathbb{E}[Y(0) \mid D=0, A=a]\Pr(D=0 \mid A=a). \tag{C.4}$$

Assumption C.4 implies that for all $\tilde{a} \leq a$,

$$\mathbb{E}[Y(0) \mid A=a] \leq \mathbb{E}[Y(0) \mid A=\tilde{a}].$$

Therefore,

$$\begin{aligned}\mathbb{E}[Y(0) \mid A=a] &\leq \inf_{\tilde{a}\leq a}\mathbb{E}[Y(0) \mid A=\tilde{a}] \\ &\leq \inf_{\tilde{a}\leq a}\Big(\Pr(D=1 \mid A=\tilde{a}) + \mathbb{E}[Y(0) \mid D=0, A=\tilde{a}]\Pr(D=0 \mid A=\tilde{a})\Big),\end{aligned}$$

where the last inequality applies equation (C.4) to all award classes $\tilde{a} \leq a$.

Substituting expression (C.3) into the left-hand side of this inequality and rearranging for

$\mathbb{E}[Y(0) \mid D=1, A=a]$ gives

$$\mathbb{E}[Y(0) \mid D=1, A=a] \leq \frac{1}{\Pr(D=1 \mid A=a)} \cdot \left[\inf_{\tilde{a} \leq a} \Big(\Pr(D=1 \mid A=\tilde{a}) + \mathbb{E}[Y(0) \mid D=0, A=\tilde{a}] \Pr(D=0 \mid A=\tilde{a}) \Big) - \mathbb{E}[Y(0) \mid D=0, A=a] \Pr(D=0 \mid A=a) \right].$$

Since $Y(0) \leq 1$, also $\mathbb{E}[Y(0) \mid D=1, A=a] \leq 1$, so taking the minimum with 1 further tightens the bound.

Lower bound on the counterfactual. Applying the lower bound on the outcome $Y(0) \geq 0$ on equation (C.3), we get $\mathbb{E}[Y(0) \mid D=1, A=a] \geq 0$. Therefore:

$$\mathbb{E}[Y(0) \mid A=a] \geq \mathbb{E}[Y(0) \mid D=0, A=a] \Pr(D=0 \mid A=a). \tag{C.5}$$

Assumption C.4 implies that for all $\tilde{a} \geq a$,

$$\mathbb{E}[Y(0) \mid A=a] \geq \mathbb{E}[Y(0) \mid A=\tilde{a}].$$

Therefore,

$$\begin{aligned} \mathbb{E}[Y(0) \mid A=a] &\geq \sup_{\tilde{a} \geq a} \mathbb{E}[Y(0) \mid A=\tilde{a}] \\ &\geq \sup_{\tilde{a} \geq a} \Big(\mathbb{E}[Y(0) \mid D=0, A=\tilde{a}] \Pr(D=0 \mid A=\tilde{a}) \Big), \end{aligned}$$

where the last inequality applies equation (C.5) to all award classes $\tilde{a} \geq a$.

Substituting expression (C.3) into the left-hand side of this inequality and rearranging for $\mathbb{E}[Y(0) \mid D=1, A=a]$ gives

$$\mathbb{E}[Y(0) \mid D=1, A=a] \geq \frac{1}{\Pr(D=1 \mid A=a)} \cdot \left[\sup_{\tilde{a} \geq a} \Big(\mathbb{E}[Y(0) \mid D=0, A=\tilde{a}] \Pr(D=0 \mid A=\tilde{a}) \Big) - \mathbb{E}[Y(0) \mid D=0, A=a] \Pr(D=0 \mid A=a) \right].$$

Since $Y(0) \geq 0$, also $\mathbb{E}[Y(0) \mid D=1, A=a] \geq 0$, so taking the maximum with 0 further tightens the bound.

Combining both inequalities yields, for each a,

$$\underline{\mathbb{E}}[Y(0) \mid D=1, A=a] \leq \mathbb{E}[Y(0) \mid D=1, A=a] \leq \overline{\mathbb{E}}[Y(0) \mid D=1, A=a].$$

Subtracting these inequalities from $\mathbb{E}[Y(1) \mid D=1, A=a]$, multiplying by $\Pr(A=a \mid D=1)$, and summing over $a \in \mathcal{A}$ (using again (C.2)) gives $\underline{\text{ATT}}_{\text{IV}} \leq \text{ATT} \leq \overline{\text{ATT}}_{\text{IV}}$. □

Relative to the unconditional Manski bounds, the Manski MIV restriction uses the award or-

dering to tighten the identified set by exploiting additional information across award categories. In particular, the upper bound on $\mathbb{E}[Y(0) \mid D = 1, A = a]$ compares award class a to all higher-potential award classes $\tilde{a} \leq a$, while the lower bound compares to all lower-potential award classes $\tilde{a} \geq a$.

Our Assumption 2 on the lower bound is similar to the Manski MIV assumption C.4 in that we also exploit an additional assumption across award categories. However, our Assumption 2 is stronger in that it restricts this across award category-comparison to specific subgroup of firms within award categories. More specifically, assumption C.4 assumes that (actual and counterfactual) outcomes are better on average across all firms in the award category, i.e., across treated and untreated firms. Instead, our assumption requires that this holds for the *untreated firms* in the better award category compared to the *treated firms* in the lower award category.

A similar comparison applies to the upper bound, but here we can compare Assumption C.4 to our *alternative* upper bound assumption 3.

D Data Appendix

This appendix documents the data collection, processing, and variable construction procedures for the analysis of Hungarian firms' exhibition participation.

In this appendix, we discuss the key steps awhile acknowledging inherent limitations of historical administrative sources.

D.1 The 1896 Budapest Exhibition

Budapest Exhibition 1896 catalog Complete official catalog of the 1896 Millennial National Exhibition containing 4,626 exhibitor records across contemporary industrial categories (agriculture excluded). Source: *Az 1896-iki Ezredéves Országos Kiállítás általános katalógusa*, edited by Mudrony Soma, Ráth Károly, and Micseh Endre (Budapest: Kosmos Ny., 1896), shelf number 282.068. The catalog comprises 21 category-specific volumes, each with independent numbering systems.

Jury Awards 1896. Official jury evaluations from the Budapest exhibition, documented in *Az 1896-iki ezredéves országos kiállítás közleményei*, issue 32, pages 1229-1236 (shelf number H 2.008 FM3/4736). Awards categorized hierarchically: Grand Certificate, Grand Medal, Exhibition Medal, and Certificate of Acknowledgment.

D.2 The 1900 Paris Exhibition and Trade Fair

Paris Exhibition 1900 catalog Hungarian-language catalog of Magyar exhibitors at the 1900 Paris International Exhibition, documenting 2,665 participants with detailed category and subcategory classifications. Source: *Az 1900. Párisi Magyar Kiállítás kalauza: Magyarország ismertetése* (Budapest: Riegler Ny., 1900), shelf number 262.540. This catalog serves as the primary treatment indicator for the empirical analysis.

D.3 Data Digitization and Text Processing

D.3.1 Digitization Protocol

All primary sources accessed in printed format underwent systematic digitization: (1) photographic capture using Hungarian National Library equipment, (2) optical character recognition via Google Document AI, and (3) systematic error correction through comparison with original printed books.

D.3.2 Text Processing Pipeline

Raw OCR output processed using Python with regular expression libraries implementing:

1. **Character encoding standardization:** Conversion of Hungarian diacritical marks and historical orthographic conventions to Unicode UTF-8 format

2. **Field extraction:** Parsing of semi-structured text entries into discrete variables using pattern recognition algorithms tailored to each source's formatting conventions

3. **Manual verification:** Systematic correction of automated extraction errors, with particular attention to firm names, addresses, and numerical data

D.4 Exhibition Participation and Awards

D.4.1 Exhibition Variables

Table D.1: Exhibition Participation Variables

Variable	Definition and Construction
`exhibition_category_YYYY`	Industry classification within exhibition framework. Original categorization preserved with harmonized coding for cross-temporal analysis.
`exhibition_subcategory_YYYY`	Detailed product-level classification within broader industry categories. Enables fine-grained analysis of exhibition participation patterns.
`exhibited_product_YYYY`	Specific products displayed at exhibition, distinct from firm's full product portfolio. Captures strategic decisions regarding exhibition representation.

D.5 Entity Resolution and Panel Construction

D.5.1 Matching Methodology

Panel dataset construction required systematic entity resolution across five independently collected datasets. The matching procedure implemented:

Pairwise Matching Algorithm:

1. **String similarity:** Jaro-Winkler distance calculation for firm names with threshold optimization balancing precision and recall
2. **Geographic consistency:** Exact and approximate location matching using standardized municipal codes
3. **Product similarity:** Jaccard similarity coefficients for product portfolio overlap
4. **Temporal consistency:** Validation of firm characteristic evolution (employment, technology) for plausibility

Correspondence Tables: Five pairwise matching tables documenting confirmed firm linkages:

- 1896 Budapest Catalog $\leftrightarrow$ 1898 Export Registry
- 1896 Budapest Catalog $\leftrightarrow$ 1900 Paris Catalog
- 1896 Budapest Catalog $\leftrightarrow$ 1904 Export Registry
- 1898 Export Registry $\leftrightarrow$ 1904 Export Registry

Master Identification Table: Comprehensive mapping consolidating direct and indirect matches to enable multi-period panel construction with both balanced and unbalanced specifications.

Table D.2: Entity Resolution Results

Dataset Pair	Total Potential Matches	Confirmed Matches	Match Rate (%)
1896 Budapest ↔ 1898 Export	1,894	1,059	55.9
1896 Budapest ↔ 1900 Paris	2,665	665	24.9
1896 Budapest ↔ 1904 Export	4,381	957	21.8
1898 Export ↔ 1904 Export	1,894	1,114	58.8

D.5.2 Match Quality and Coverage

D.6 Sample Characteristics and Descriptive Statistics

D.6.1 Dataset Coverage and Temporal Scope

Table D.3: Dataset Characteristics

Dataset	Observations	Time Period	Processing Status	Key Innovation
Budapest 1896	4,626	1896	Complete*	Universal exhibition census
Export 1898	1,894	1897-1898	Complete	Export-ready classification
Paris 1900	2,665	1900	Complete	International treatment
Export 1904	4,381	1904-1906	Complete	Comprehensive registry

*Excludes education and agriculture categories

D.7 Measurement Limitations and Econometric Considerations

D.7.1 Historical Source Constraints

Measurement Error: Firm-reported employment and capital measures subject to strategic reporting incentives and standardization challenges across heterogeneous enterprises.

Survival Bias: Panel analysis necessarily conditions on firm survival through the observation period, potentially understating treatment effects for marginal firms.

D.7.2 Entity Resolution Limitations

Several issues arose during the process that may generate both types of errors. We looked at matches manually but it remains possible.

Matching Precision: Name variations, mergers, and reorganizations

Geographic Mobility: Firm relocation between observations may complicate identification. (A systematic address tracking suggests limited geographic mobility for manufacturing enterprises.)

D.7.3 Variable Construction Issues

Industry Classification Harmonization: Evolving industrial classification systems across datasets necessitate concordance tables with potential misclassification of diversified firms.

Temporal Aggregation: Registry publication dates may not correspond precisely to data collection periods, introducing timing measurement error requiring robustness checks with alternative temporal specifications.

E Exhibition Rules and Institutional Framework

The credibility of our empirical strategy relies on understanding the institutional rules governing both exhibitions. The Budapest exhibition provided an objective measure of export potential through its independent jury system, while the Paris exhibition's selection constraints created the variation necessary for identification.

E.1 Budapest 1896 Millennial Exhibition Rules

The 1896 Budapest exhibition was designed as a comprehensive showcase of Hungarian industry and served as a "trial run" for identifying firms with export potential. The exhibition attracted over 9,000 exhibitors across all sectors, with mandatory evaluation of all commercial exhibits by an independent jury system.

E.1.1 Jury Independence and Structure

The jury operated independently from any future selection processes. According to the official regulations, jury members were selected through a mixed process: half elected by exhibitors within each industry group, half appointed by the exhibition committee in consultation with relevant ministries. The jury consisted of 480 members across 34 specialized industry groups, with each group electing its own president, vice-president, and rapporteur by secret ballot and simple majority.

Crucially, jury members and their associated firms were explicitly excluded from receiving awards in their respective groups, ensuring objective evaluation. Experts serving in advisory roles were similarly excluded from awards in the specific groups where they provided assistance. State authorities, institutions, and enterprises were also ineligible for awards, with their items marked with appropriate labels.

E.1.2 Award Categories and Hierarchy

The jury awarded four distinct categories of recognition, listed below in descending order of prestige. Our labels used in the main text are given in parentheses.

1. **Honorary Diploma (Gold)**: The highest award, given to exhibitors who had earned "national merit" through their contributions, awarded by the jury council based on group jury proposals. A maximum of 100 Honorary Diplomas could be granted. Since some exhibits were jointly produced by several firms, the 100 diplomas correspond to 135 unique firms in our sample.

2. **Millennial Grand Medal (Silver)**: Awarded to exhibitors who "demonstrated exceptional advancement in their field, elevated their industry, introduced new types of production or methods, contributed to mass production or exports, and enhanced the prestige of the exhibition through the excellence of their exhibits."

3. **Exhibition Medal (Bronze)**: Awarded to exhibitors who demonstrated merit worthy of distinction.

4. **Certificate of Recognition (Honorable Mention)**: Granted to exhibitors not eligible for higher awards but whose work deserved acknowledgment.

Except for honorary diplomas, there was no limit to the number of medals and certificates awarded. Each medal was accompanied by a certificate listing the specific merits upon which the award was based.

E.1.3 Award Criteria

The regulations often explicitly specified export capability as a primary evaluation criterion. Common citations in the jury's stated rationales for awards included: "good/excellent export capability" (kivitelképesség), "competitiveness" (versenyképesség), "introducing a new industry in Hungary (új iparág bevezetése)," "good/excellent work/effort (jó/kiváló munka), or performance (teljesítmény)," and "good/excellent progress or innovation (jó/kiváló haladás)." The emphasis on export potential was central to the exhibition's role as a trial run for identifying firms suitable for international representation.

E.1.4 Evaluation Process

All exhibited items underwent mandatory evaluation unless the exhibitor explicitly declined within a set deadline. The jury conducted systematic on-site inspections and decided for each exhibitor: (1) whether they should receive a distinction and why, (2) which award should be granted, and (3) what specific citation should be used.

Group jury decisions were made by majority vote, with the president's vote deciding in case of ties. When decisions were not unanimous, minority opinions could be recorded upon request. The evaluation process was conducted under strict confidentiality, with no appeals permitted against jury council decisions.

E.2 Paris 1900 Exhibition Rules and Selection Process

The Paris exhibition operated under a sophisticated three-tier international jury system, but Hungarian participation was determined by a separate domestic selection process subject to multiple constraints beyond export potential.

E.2.1 Scale and Commercial Purpose

The 1900 Paris Exposition Universelle was explicitly designed as the largest international commercial gathering to date, featuring 76,000 exhibitors from 56 countries across 121 product subcategories. The exhibition's commercial objectives were formally recognized, with contemporary documentation describing it as providing "favorable opportunities for establishing commercial connections" and facilitating face-to-face meetings between international buyers and sellers.

E.2.2 Hungarian Selection Constraints

While the primary goal was promoting exports by selecting high-potential firms, internal government records reveal three additional constraints that prevented perfect sorting based solely on export potential:

1. **Space Limitations**: Limited pavilion space in Paris restricted total participation numbers.
2. **Representational Balance**: The selection committee sought balanced representation across industries and regions to present a comprehensive picture of Hungarian production capabilities.

3. **Transport Costs**: High costs for shipping exhibits, particularly heavy machinery and bulky products, influenced selection decisions and created logistical constraints.

E.2.3 Product Category Organization

Exhibitors were classified into 18 broad categories and 121 subcategories based on their specific exhibits rather than their full product portfolios. This classification system determined competitive exposure at the exhibition, with subcategory sizes varying from hundreds to thousands of exhibitors. Categories ranged from narrowly defined products (e.g., "shoes") to broader industrial classifications (e.g., "chemicals").

E.2.4 Commercial Networking Framework

The exhibition operated from April to November 1900, providing extended opportunities for commercial relationship formation. The 48 million visitors included merchants, industrialists, and potential customers, creating an unprecedented concentration of international commercial contacts. Contemporary accounts emphasize active commercial negotiations occurring alongside product displays throughout the exhibition grounds.

These institutional details underscore that Budapest awards provide a credible measure of export potential determined by independent expert evaluation using explicit export-focused criteria, while Paris selection involved strategic balancing across multiple objectives, creating the selection patterns we exploit for identification.