

Institutional Narratives in the Art Market

Evidence from Museum Leadership Transitions

Research Question

Does exogenous variation in institutional attention to art categories predict subsequent movements in category-level auction prices? This paper exploits the staggered turnover of directors and chief curators at major museums as a natural experiment in the reallocation of cultural legitimacy across art categories.

Motivation and Core Logic

The determinants of art prices remain poorly understood precisely because artworks lack the cash flows, use value, or replacement cost that anchor valuation in other asset classes. In their absence, market participants must rely on socially produced assessments of quality and importance, assessments in which museums, curators, and exhibition programs play a constitutive rather than merely reflective role. The art historical literature since Becker and Moulin has emphasized that cultural value is not a fixed property of objects but an outcome of institutional processes: selection for exhibition, inclusion in permanent collections, placement within art historical narratives. If these processes carry informational content for collectors and dealers, then shocks to institutional programming should propagate into transaction prices. Yet the causal relationship between institutional legitimation and market valuation has proven difficult to isolate empirically, because institutions, markets, and broader cultural taste tend to co-move.

This paper exploits museum leadership transitions as a source of plausibly exogenous variation in institutional attention to resolve precisely this identification problem. When a new director assumes control of a major institution, they reliably redirect exhibition programming, acquisitions, and publication activity toward categories aligned with their prior curatorial identity. This reallocation is driven by career formation, intellectual commitment, and institutional strategy, not by speculative beliefs about which art categories will appreciate at auction. The appointment itself is determined by board governance dynamics, donor politics, and candidate

availability. Neither the timing nor the categorical direction of the programming shift is plausibly caused by contemporaneous auction price trends. The contribution, accordingly, is not to claim that institutional legitimation is the sole or even primary determinant of art prices, but rather to isolate one channel among many and to test whether it carries a detectable, if partial, causal effect in a market where multiple value registers operate simultaneously and often in tension with one another.

Empirical Strategy

The design proceeds in two stages that decompose the causal chain.

Stage one (first stage): leadership turnover shifts institutional attention.

I identify approximately 30 to 50 major leadership transitions at 15 to 25 significant museums in North America and Europe between 1975 and 2020. For each incoming leader, I code their categorical orientation ex ante using observable career history: prior curatorial positions, exhibition record, publication topics, and stated programmatic vision at the time of appointment. Categories aligned with the incoming director are designated as “treatment” categories at that institution; categories aligned with the departing director or orthogonal to both serve as controls. I then estimate whether treatment categories experience a differential increase in exhibition frequency, catalog production, and acquisition activity at the affected museum relative to control categories, using a staggered difference in differences framework. This stage should be close to mechanical and serves primarily to validate the instrument. For demonstrative purposes, the first-stage specification takes the form:

$$A_{cmt} = \alpha + \beta(Aligned_{cm} \times Post_{mt}) + \gamma_c + \delta_t + \mu_m + \varepsilon_{cmt} \quad (1)$$

where A_{cmt} denotes the institutional attention measure (exhibition share, catalog entries, or acquisitions) for art category c at museum m in year t ; $Aligned_{cm}$ is an indicator equal to one when category c matches the incoming director’s curatorial orientation at museum m ; $Post_{mt}$ equals one for years following the leadership transition at museum m ; and γ_c , δ_t , and μ_m absorb category, year, and institution fixed effects, respectively. The coefficient β captures the differential change in institutional attention for aligned categories after the transition.

Stage two (reduced form): do price indices respond?

Using auction transaction data from Artnet, Blouin, or comparable sources, I construct category-level price indices (hedonic or repeat sales) at annual or semi-annual frequency. The core specification tests whether treatment categories at a given transition event experience differential price growth relative to control categories in the three to seven years following the leadership change, after absorbing category fixed effects, time fixed effects, and institution fixed effects. The staggered timing of transitions across institutions and the cross-sectional variation in categorical direction together provide identification. For demonstrative purposes, the reduced-form price specification takes the form:

$$\ln P_{ct} = \alpha + \sum_m \beta (\text{Aligned}_{cm} \times \text{Post}_{mt}) + \phi \mathbf{F}_t + \psi \text{Trend}_{ct} + \gamma_c + \delta_t + \varepsilon_{ct} \quad (2)$$

where $\ln P_{ct}$ is the log category-level price index for art category c in year t ; the summation aggregates treatment exposure across all museum transitions m to which category c is subject; \mathbf{F}_t is a vector of broader financial market controls (equity returns, interest rates, aggregate art market indices) included to strip out co-movement with asset prices unrelated to institutional legitimation; Trend_{ct} captures category-specific pre-transition price trends to purge idiosyncratic momentum that might otherwise pollute the estimate; and γ_c and δ_t are category and year fixed effects. The coefficient β identifies the differential price response attributable to the reallocation of institutional attention, net of both macrofinancial conditions and pre-existing category-level dynamics.

A clarification on panel structure and the nature of the identifying variation is warranted. The unit of observation in these regressions is not the leadership transition itself but rather category-by-museum-by-year in the first stage and category-by-year in the reduced form. Each transition event simultaneously treats multiple art categories, and each category in turn contains potentially hundreds or thousands of auction lots feeding into the price index. Even with 15 to 20 clean transition events, the resulting panel is therefore large in observational terms. The cross-sectional variation in *which* categories are treated at each transition, as determined by the incoming director's curatorial identity, is doing substantial work: it is this multiplier, not the count of transitions per se, that generates the dimensionality of the estimation sample. What the design is estimating, in essence, is the price effect of staggered breaks in institutional attention, where each break redirects legitimation toward a different subset of categories at a different point in time.

Where the limited number of independent transition events does bite is inference. If standard errors are clustered at the transition level, which is arguably the correct level since treatment assignment originates there, then 15 to 20 clusters places the analysis squarely in the “few clusters” regime where cluster-robust variance estimators are downward biased. Reliable p -values will therefore require wild cluster bootstrap or randomization inference procedures. Moreover, in light of recent advances in the staggered difference in differences literature (Sun and Abraham, 2021; Callaway and Sant’Anna, 2021), having few treatment cohorts with heterogeneous timing makes the decomposition of the average treatment effect more fragile, and I will report cohort-specific estimates alongside any aggregated coefficient. The concern, in short, is not insufficient data for estimation but rather insufficient independent sources of identifying variation for robust inference, and it is this distinction that disciplines how results should be interpreted and reported.

I additionally test for dosage effects along two dimensions: institutional prestige (transitions at MoMA or the Tate should matter more than transitions at smaller institutions) and magnitude of programming shift (larger reallocations of exhibition share should produce larger price responses). Both predictions are derived directly from an information diffusion account and would be difficult to rationalize under pure reverse causality.

Data Requirements

Three datasets, all feasible to construct. First, a leadership transition database: director and chief curator appointments at major museums, sourced from institutional archives, press coverage, and existing museum history scholarship. Second, categorical orientation coding: the incoming leader’s prior career record, coded into predefined art categories (movement, period, medium, and geographic origin). This must be done using information available strictly before the appointment date. Third, exhibition and acquisition records at affected institutions, which are available through museum annual reports, published exhibition histories, and digital catalog archives. Fourth, category-level auction price data, available commercially.

Identification Threats and Responses

The primary concern is that boards appoint directors in order to pivot toward categories the board believes will appreciate. This is implausible at the level of major public museums, where board selection criteria center on fundraising capacity, public profile, donor relationships, and institutional vision rather than speculative art market positioning. I will test this directly by examining whether treatment categories show any differential price trend in the years *before* the leadership transition. Pre-trend divergence would be evidence of selection; its absence supports the exogeneity claim.

A secondary concern is that prominent directors may be appointed *because* certain categories are already gaining cultural momentum, and both the appointment and subsequent prices reflect this momentum. I address this by controlling for pre-transition institutional attention trends and by examining cases where the incoming director's orientation represents a sharp break from the departing leader's program, which isolates reallocation from continuation.

Expected Contribution

If institutional attention shocks predict price movements, this provides direct evidence that qualitative, socially produced information enters prices in markets characterized by extreme heterogeneity and contested value. If the relationship is null, that is equally informative: it would suggest that auction markets operate largely independently of institutional legitimation, which challenges a foundational assumption in both cultural economics and the sociology of art. Either result speaks to the broader question of how prices form when fundamental value cannot be derived from observable cash flows or use value, a question relevant well beyond art markets.