

Research Statement

Samiha B. Tariq

PhD Candidate in Economics, Southern Illinois University Carbondale

samiha.bintetariq@siu.edu

I study how prices form in markets that standard empirical designs handle badly. The markets I work in share a feature: the obvious way to answer the question gives an answer that does not hold. In each project the substantive finding has come after establishing what the data can actually identify, and in more than one case the second step overturned the first. That sequence has become the method rather than an accident of it, and it is what connects work on digital collectibles, fine art at auction, household credit and marriage decisions.

Job market paper

Pixels, Scarcity, or Brand? What Visual Features Are Worth in NFT Markets asks how much of what people pay for art is payment for how it looks. Cultural economics has posed this question of paintings since Rosen, and has always been limited by the same obstacles: each painting sells once or twice, its attributes are entangled with its maker, and the maker is entangled with everything else. Non-fungible tokens appear to remove all three. Every pixel is recorded rather than described by an expert, provenance is complete and public, and the same image sells repeatedly.

I measure 196 computer-vision descriptors on 23,997 images and match them to 93,701 sales across 26 Ethereum collections. The design then takes seriously three properties of generative collections that a token-level hedonic must survive: visual style is largely a collection-level attribute, the effective unit of inference is the collection rather than the sale, and scarcity is a deterministic function of the same attribute layers that generate the image.

Taking those seriously reverses the standard result. Identifying within collection, seven of thirteen leading coefficients change sign. Conducting inference at the collection level, where there are 26 independent units rather than 93,701, one descriptor with an analytic t -statistic of +9.32 returns a bootstrap p -value of 0.246; after correcting for 196 simultaneous tests, no descriptor carries an identified premium under any standard multiplicity correction. Scarcity does, at 0.154 log points per within-collection standard deviation, and is the only regressor that clears inference at that level.

The paper also proves why the prevailing design reports premia that are not there. With token-constant regressors and a token random intercept, the estimated random effects satisfy $X'\hat{u} = 0$ exactly, for any variance ratio, so the intercept is forced orthogonal to the included regressors and cannot absorb correlated omitted variation. In a simulation where a feature's true effect is exactly zero, that design returns +3.03. The result is not specific to digital assets: it applies to any hedonic design with object-level random intercepts and object-constant characteristics, which describes much of the empirical literature on art, housing and consumer products.

Two further results give the paper its positive content. Pixels predict within-collection scarcity out of sample, non-linearly, through texture and busyness rather than composition, so rarer tokens simply contain more visual information. And a visual index built to predict scarcity is priced, surviving both the measured scarcity score and a flexible version of it, until the full attribute vector enters, at which point it falls to a fifth of its size and reverses sign. Visual features price these tokens to the extent that they reveal which attributes a token carries, and no further. The null is not a power problem: eleven times as many tokens

would move nothing, because what limits precision is dispersion of coefficients across collections rather than noise within them.

Monetary transmission into cultural assets

The Fed at the Auction House: A Monetary Transmission Puzzle in Old Master Art, with Zsolt Becsi, asks whether monetary policy surprises reach an asset that pays no cash flow, exists in fixed supply, and is bought mostly by the wealthy. Using a new lot-level panel of 478,199 Christie's auction records matched to four published series of FOMC policy surprises, we find that they do, with the wrong sign. A hawkish surprise raises the price an Old Master lot fetches relative to its pre-sale estimate, and the effect is largest in New York, where dollar bidders compete most directly. The same surprises move seven benchmark financial assets in their textbook directions, leaving Old Master art the lone asset with the wrong sign. Because the response is identified from thirty FOMC-coinciding sale dates, we lead with a wild cluster bootstrap rather than the cluster-robust standard error, and report a permutation test alongside it, so that the confidence attaching to the estimate is visible as the inference procedure becomes more demanding. We read the evidence as a store-of-value channel of monetary transmission.

What does move these prices

If appearance is not separately priced, something else is doing the work. ***Hype Has Worth: Attention, Sentiment, and NFT Valuation in Major Ethereum Collections*** merges 87,696 secondary-market sales with Reddit discourse measures across 25 collections, absorbing visual heterogeneity through a transparent within-collection index and separating persistent cross-collection differences from within-collection variation through a Mundlak decomposition. Valuations align most strongly with sustained collection-level attention rather than short-horizon noise, and attention is most informative measured as cumulative engagement across several prior windows. Read with the job market paper, the pair says that what is priced in these markets is scarcity and sustained attention, not appearance.

The same problem outside finance

Smiles Before Vows? Positive Affect and the Marriage Decision, with Zsolt Becsi, shows the measurement problem is not confined to asset markets. Studies of what moves partnership test against the share of adults currently in a union, a stock that contains people who formed unions decades ago and whom nothing measured today could have moved. Splitting the population by whether the decision is still ahead, which a census permits because it records age, a standard deviation of national Positive Affect raises the partnered share of cohorts at the ages of union formation by 1.94 points relative to those past them, across 109 countries. Pooled across all ages the same coefficient cannot be distinguished from zero, which is why two literatures testing against the stock found nothing.

Published work

My co-authored work studies how global shocks propagate across markets. *Unravelling Determinants and Interconnectedness among Global Risks and Financial Stress in South and Southeast Asia* (*International Review of Economics & Finance*, 2026) maps how geopolitical risk, oil and gold market uncertainty and global conditions transmit into eight Asian economies, and identifies which markets absorb shocks and which send them. *Time-Varying Connectedness Among Oil Price Shocks, Global Conditions, and Financial Stress in South and Southeast Asian Markets* (*Computational Economics*, 2025) shows that these linkages shift with frequency and crisis episode, with long-run spillovers the more persistent.

Agenda

The job market paper identifies its own binding constraint, and the next project follows directly from it. What limits precision is the number of collections, not the number of tokens, so a few hundred tokens drawn from each of a hundred collections would settle whether any common visual premium exists at all, a question the present design cannot answer either way. That study is also cheaper than the one I have already run.

Beyond it, I want to extend the identification result. The proposition applies wherever objects carry time-invariant characteristics and the design absorbs object-level heterogeneity with a random intercept, which covers a substantial body of work in art, real estate and product hedonics. Quantifying how much of the reported premia in those literatures is attributable to the specification rather than to the characteristics is a tractable and, I think, useful project.

I am also interested in extending the auction work: whether the store-of-value response we find in Old Masters appears in other non-yielding collectibles, and whether it strengthens where currency exposure is more concentrated. Across all of it, my aim is the same. I would rather establish a narrow result that holds than a broad one that does not, and I have found that working out why the broad result failed is usually where the contribution turns out to be.