

Qualities or Inequalities?: How Gender Shapes Value in the Market  
for Contemporary Art

by

Taylor Whitten Brown

Department of Sociology  
Duke University

Date: \_\_\_\_\_

Approved:

\_\_\_\_\_  
Christopher A. Bail, Advisor

\_\_\_\_\_  
James Moody

\_\_\_\_\_  
Mor Naaman

\_\_\_\_\_  
Martin Ruef

\_\_\_\_\_  
Lynn Smith-Lovin

\_\_\_\_\_  
Stephen Vaisey

Dissertation submitted in partial fulfillment of  
the requirements for the degree of Doctor  
of Philosophy in the Department of  
Sociology in the Graduate School  
of Duke University

2021

ABSTRACT

Qualities or Inequalities?: How Gender Shapes Value in the Market  
for Contemporary Art

by

Taylor Whitten Brown

Department of Sociology  
Duke University

Date: \_\_\_\_\_

Approved:

\_\_\_\_\_  
Christopher A. Bail, Advisor

\_\_\_\_\_  
James Moody,

\_\_\_\_\_  
Mor Naaman

\_\_\_\_\_  
Martin Ruef

\_\_\_\_\_  
Lynn Smith-Lovin

\_\_\_\_\_  
Stephen Vaisey

Dissertation submitted in partial fulfillment of  
the requirements for the degree of Doctor  
of Philosophy in the Department of  
Sociology in the Graduate School  
of Duke University

2021

Copyright by  
Taylor Whitten Brown  
2021

## **Abstract**

How does gender inequality persist in the art world today? Or, more generally, what role do social characteristics like gender play in markets for cultural goods, such as art? That is the focus of this research. Using a novel dataset of 255,887 contemporary artworks produced by 18,624 artists and gleaned from an online marketplace, I employ the case of gender in the art world to investigate how social characteristics of producers can impact market outcomes and structures. Although there is prominent scholarship on product markets and inequality within sociology, questions such as these are rarely posed. Work generally focuses on the quality of goods and on the status of producing organizations, without attention to individual producer characteristics, including gender.

The first study of this dissertation implements machine learning classification to examine whether female and male artists produce artworks with different characteristics. These analyses rely on a taxonomy of over 1,000 art-relevant features, coded by a team of art historians, to describe the disciplines, physical attributes, styles and periods, object types, and settings of each artwork in the dataset. I find that artworks by women and men do not substantively differ on the majority of aesthetic, conceptual, or material features that they depict. While some, less common, features of

art appear more in work by women or men, by in large these two groups of producers do not bring different products to the art market.

Studies two and three of this dissertation move to address alternative hypotheses for disparity in the economic outcomes of women and men in the contemporary art market. With mixed effect regression, I test whether artworks by women are priced differently than artworks by men, even after accounting for the categories and features they depict. I find that art by women is listed at a discount of approximately 10 to 12 percent relative to art by men. I also find that, of those art qualities that differ in use between women and men, qualities of art predominantly made by women are valued less than those predominantly made by men, net of who creates them. In combination, these findings echo and extend calls to value the labor of women and men comparably. They also broaden our understanding of the potential for social status characteristics, like gender, to act as organizing structures in the production, meaning, and valuation of markets.

## **Dedication**

To femininity, artists, and the will of women.

# Contents

Abstract.....	iv
List of Tables .....	x
List of Figures .....	xi
Acknowledgements .....	xii
1. Introduction .....	1
2. Case: the contemporary art market .....	9
3. Study A .....	14
3.1 Product markets, categories, and features .....	14
3.2 Product markets and status.....	19
3.3 Labor markets and status .....	22
3.4 Hypotheses .....	26
3.5 Part 1: Do women make art with different characteristics than men?.....	27
3.5.1 Data.....	28
3.5.2 Methods .....	30
3.5.3 Results .....	34
3.5.3.1 Artwork-level random forest classification.....	34
3.5.3.2 Artist-level descriptive statistics .....	37
3.5.4 Discussion.....	41
3.6 Part 2: Is artwork by women valued less than artwork by men? .....	44
3.6.1 Data.....	44

3.6.2 Methods .....	46
3.6.2.1 Log linear mixed effect regression.....	46
3.6.2.2 Covariate balancing propensity score weights.....	47
3.6.3 Results .....	51
3.6.3.1 Descriptive statistics .....	51
3.6.3.2 Log linear mixed effect regression analysis .....	53
6.2.4 Discussion.....	58
4. Study B.....	66
4.1 Material and economic evidence of devaluation theory.....	67
4.2 Symbolic evidence of devaluation theory .....	70
4.3 Devaluation theory in the context of product markets.....	72
4.4 Hypotheses .....	76
4.5 Data and method .....	77
4.6 Results .....	80
4.7 Conclusion .....	85
4.8 Discussion.....	87
5. Conclusion.....	92
Appendix A.....	100
A.1 Art Genome Project vis-à-vis academic taxonomies of art.....	100
A.2 Strengths and Limitations of the Art Genome Project .....	100
Appendix B .....	105
B.1 Selecting artworks for analysis.....	105



B.2 Imputation of missing gender data .....	106
Appendix C .....	108
C.1 Feature selection and modification for random forest classification .....	108
C.2 Classifier robustness checks .....	109
Appendix D .....	111
D.1 Mixed effect log linear model of gallery listing price for artwork (Model 2), by medium .....	111
Appendix D.2 Robustness check of mixed effect log linear model of gallery listing price for artwork (Model 4) .....	112
Appendix E. Future work on the network structure of the contemporary art market .....	114
References .....	122
Biography .....	141

## List of Tables

Table 1: Number of artworks ascribed to each feature.....	33
Table 2: Number of features ascribed to each artwork.....	33
Table 3: Descriptive statistics for variables used in regression analysis .....	51
Table 4: Zero-order correlation of variables used in regression analysis .....	52
Table 5: Mixed effect log linear model of gallery listing price for artworks .....	55
Table 6: Quantile regression of log gallery listing price for artworks .....	56
Table 7: Share of market value and volume, by quintiles of price and gender.....	57
Table 8: Count and share of artworks, by quintiles of listing price and gender.....	58
Table 9: Count and share of artists, by quintiles of listing price and gender .....	58
Table 10: Mixed effect log linear model predicting gallery listing price for artworks.....	82
Table 11: Mixed effect log linear model of gallery listing price for artwork, by medium (Model 2 in manuscript).....	111
Table 12. Artist network variables derived from the network of artists tied by mutual gallery representation.....	119

## List of Figures

Figure 1: Confusion matrix and precision-recall plot of results from random forest classification .....	36
Figure 2: Robustness checks of random forest classification. ....	37
Figure 3: Logarithm of relative frequency ratios for artist features used by women and men, by overall feature frequency .....	40
Figure 4: Plot of standardized mean difference in artistic features for artwork by women and men, adjusted and unadjusted by covariate balancing propensity score weights ....	50
Figure 5: Artworks plotted by the gender-typicality index, and by the gender of the artist (n = 102,908).....	80
Figure 6: Predicted listing price for artworks, derived from Model 4.....	84
Figure 7. Network representation of art galleries tied by mutual representation of contemporary artists .....	117
Figure 8. Network representation of contemporary artists tied by mutual representation by at least one gallery .....	118
Figure 9. Network representation of contemporary artists tied by the features represented in their artwork.....	121

## Acknowledgements

Foremost, my committee: Christopher Bail, James Moody, Mor Naaman, Martin Ruef, Lynn Smith-Lovin, and Stephen Vaisey. Especially Chris, for inspiring my first steps into computational social science, and for providing me with incomparable opportunities to make an intellectual home there. And Jim, for the firehose of your Social Networks class and your keen respect for data visualizations, which started this project. Martin and Mor, the results of this work would not be as accurate without your insights. Lynn, your class on Relational Sociology will forever shape my imagination. And Steve, thank you for your support as graduate director, statistical advisor, and friend. More broadly, this dissertation has been molded by the sociology departments of both UNC-Chapel Hill and Duke. Their influence will now go wherever I do.

Data from Artsy.net was indispensable to this dissertation. I am indebted to the company, and specifically to Maddie, Rachel, and the Art Genome team of historians for their dedication to the science of taxonomizing art. Also, Orta, who shared data in the spirit of open science, and Jamis, for assisting me with the API.

My friends and fellow graduate students at UNC and Duke have more than once been my anchors to well-being and insight. Ali, Josh, Lauren, LesLeigh, Freido, Marcus, Mary Beth, Steven, and others. Furthermore, I am immensely grateful to the younger students, with whom I bonded throughout the pandemic and by whom I've been inspired in our efforts make the discipline more equitable. I am grateful for the

community gained via the Summer Institute in Computational Social Science. Abdullah, Allison, Chris, Ian, Josh, Kat, Nicolò, Tina, Raquel, Ridhi, Xiao, and so many others. Every day I am inspired by at least one of you in method and goodness. Likewise, my team on Core Data Science at Facebook, whose skill in science is imperative. Winter Mason, thank you for accepting, pushing, and consistently mentoring me.

This research was aided by every coffeeshop I worked from, and by Peter Bearman and Paul di Maggio, who invited me to visit at Columbia and New York University. Graciously, you provided me with space to study and a sense of community in my new home. So too every administrator who navigated institutions on my behalf—especially Jessica—your contribution is not absent in my gratitude.

While it was always my intention to study how culture perpetuates gender inequality, I did not set out to study the art world. That said, in doing so I have gained a conviction to the essential role taken by artists in society, and to the effort of encouraging a diversity of perspectives to thrive in their profession. I therefore acknowledge the anonymous art group Guerrilla Girls, who collected many of the first statistics on gender inequality in the art world, one of which provided my entre to the issues of power at play. And Clare McAndrew, who produces essential reports.

Finally, indelibly, I acknowledge Spencer. The worth of your love and patience to this endeavor cannot be measured. Thank you, so much.

# 1. Introduction

Only five percent of artworks on display in modern and contemporary sections of the Metropolitan Museum in New York City were by women in 1989. In 2005, that figure dropped to 3 percent, and in 2012 it stuttered up to only 4 percent, with other prominent institutions conveying similar statistics. As of 2018, only 13 percent of artworks in major US institutions and 33 percent of artworks in global exhibitions were by women (*The Art Market 2019*, 2019; Topaz et al., 2018). No woman has entered the top one hundred auction sales for paintings, and yet again in 2020 no woman entered the top ten. Women comprise only 35 percent of artists represented by galleries as of 2020 (down from 39 in 2019), despite comprising the majority of graduates with fine arts degrees (*The Art Market 2020*, 2020; Townsend, 2017).<sup>1,2</sup>

How does such gender inequality persist in the art world today? Or, more generally, what role do social characteristics like gender play in markets for cultural goods, such as art? That is the focus of this research. Using a large dataset of 255,887 contemporary artworks drawn from an online marketplace, I employ the case of gender in the art world to investigate how social characteristics of producers can impact market outcomes. Although there is prominent scholarship on product markets and inequality

---

<sup>1</sup> This statistic reflects artists represented by only one gallery. For artists represented by more galleries—a state often associated with career success—the rate of female representation is lower.

<sup>2</sup> This statistic is specific to the United States. In 2015, women in the United States earned 70% of Bachelors of Fine Arts and 65–75% of Masters of Fine Arts degrees (Townsend, 2017).

within sociology, questions such as these are rarely posed. Work generally focuses on the quality of goods and the status of producing organizations, without attention to individual producer characteristics, including gender (Fligstein & Dauter, 2007; Fourcade, 2007).

As an unexpected omission, the groundwork for research into individual producer characteristics has nevertheless been laid. There is a close intellectual proximity between the product market literature and the labor market literature in sociology, the latter of which has developed great veins of stratification scholarship to scrutinize disparities on the basis of worker social demographic characteristics (G. S. Becker, 1957; Du Bois, 1899; England, 2005). The labor market literature is thus well positioned to draw together cultural and economic sociology for a better understanding of how the economic valuation of products may relate to the social valuation of their producers.

In the specific context of cultural product markets, such as art, where goods are both commodities and symbolic objects (Bourdieu, 1983), one may be surprised to acknowledge an absence of research on the characteristics of producers. After all, producers in these markets are generally expected to create work that in some way aligns with their identity. Here, however, the absence of sociological research on producer characteristics is somewhat comprehensible, if only because academic trends have been explicit. For instance, while social theory on consumption has long

emphasized that group identities structure and are structured by the patterns of goods consumed by group members (Bourdieu, 1979; DiMaggio, 1987; Simmel, 1957; Veblen, 1899), it has also persuasively argued that the careers, market outcomes, and even the very goods that producers create are collective endeavors, not just the result of individual artists whose genius or identity must be scrutinized (de la Fuente, 2007). Placing the study of cultural production squarely in contrast to the approach taken by art historians, therefore, modern sociologists have tended to inquire less about individual producers, and more about the social coordination and deliberation of the many actors who surround them (H. S. Becker, 1974, 1982; Padgett & Ansell, 1993; H. C. White & White, 1993). Phrased by Bourdieu, “the question is not what the artist creates, but who creates the artist” (Bourdieu, 1993, p. 147).

This intellectual perspective—foregrounding the society that surrounds an artist over the artist themselves—has yielded many important insights. It has provided tools to demystify the supposed plethora of auteur, white, male artists (Nochlin, 1971), exposed the role of networks in artistic careers and historical legacies (Giuffre, 1999; Lang & Lang, 1988; Padgett & Ansell, 1993), and generated alternative understandings for the evolution of artistic movements and markets (H. C. White & White, 1993). Yet in focusing on all that surrounds the artist and their work, notice of the relationship between the identity of the artist and the qualities of their work has been obscured,



sidelining scrutiny of whether and how this relationship carries social meaning that may influence market behaviors and outcomes.

In the present dissertation, I argue that economic, gender, and cultural sociology have much to gain from one other in the analysis of product markets and social inequality. I use their intersection to address the case of gender inequality in the contemporary art market, and in so doing attend to details both of producers and of their products. The crux of the evidence provided by my analyses is that individual producer characteristics—in this case, gender—can become salient status characteristics that shape markets via patterns of production, outcomes of value, and notions of quality. As a result, these identity characteristics entwine markets with the broader social world, making it difficult to extricate and imperative to address their presence in social scientific research.

This project provides unique empirical analyses for both new and longstanding hypotheses on product markets and the art world. Through three studies, I show that the relationship between gender and the economic outcomes of contemporary artists (i.e. artists working today, or in recent decades) is an artifact of both differences in the qualities of work produced by women and men, and biases in the value ascribed to that work through labor and meaning. Yet while both are at play, bias in valuation better accounts for the economic gap between sexes.

In my first set of analyses, I draw on a dataset of 255,887 recent artworks produced by 18,624 artists and gleaned from an online marketplace. Each work is described by a taxonomy of over 1,000 art-relevant features, coded by a team of art historians. The features relate to the disciplines, physical attributes, styles and periods, object types, and settings of each work. Using this data, I implement machine learning classification to examine whether female and male artists produce artworks with different characteristics. If they do, this may account for some degree of imbalance in market success. I find, however, that artworks by women and men do not substantively differ on the majority of aesthetic, conceptual, and material features that they depict. While some, less common, features of art appear more in work by one gender of artist or the other, by in large these two groups of producers do not bring qualitatively different products to the art market.

Building on these results, my second set of analyses address hypotheses for disparity in the economic outcomes of female and male artists today. With mixed-effect regression, I test whether artworks by women are priced differently than artworks by men, even after accounting for the categories and features of their work. I find that women's art is listed at approximately a 10 percent discount to men's. This is comparable to the gender wage gap in the labor force at large, which persists at between 8.5 and 20.7% in the US, depending on population and statistical adjustment (F. D. Blau & Kahn, 2016). This also supports the claim that gender is a salient status characteristic

in the contemporary art market, and that neglecting its role in shaping outcomes of valuation will inhibit our ability to understand the market or address its inequalities.

I do not simply aim to press an argument of bias in judgment against female artists, however—true and relevant though that may be. I also aim to explore whether categories and features of art constitute an additional layer of inequality and a stickiness to the status quo that has historically benefitted men. The assumption here is that gender beliefs ascribed to producers (artists) may link with the aesthetic, conceptual, and material qualities of their products (artworks), and thereby lead logics of valuation to be biased against qualities of art typically associated with women—even when those qualities are produced in the work of a man. As the third set of analysis in this manuscript establish, artistic qualities predominantly associated with the work of women are devaluated relative to qualities predominantly associated with the work men. This goes beyond the devaluation of women’s labor and broadens our understanding of the potential for social status characteristics, like gender, to act as an organizing structure of product markets.

I charge, based on the findings of this dissertation, that gender equality in the art market today will remain out of reach if the only change is to value the labor of women and men equally. There must also be an equal valuation of comparable aesthetic, conceptual, and material features of art that these two populations produce. This is an

argument often repeated in studies of the labor market, but now aptly echoed in a product market.

This dissertation makes several contributions to the study of gender, culture, and markets in sociology. First, I approach the study of product markets with the aim of analytically distinguishing between the compositional elements of inequality, which are due to differences in the goods made by producers, and valuation elements, which are due to differences in judgement on what is made. Second, I rely on several analytical methods, including machine learning and propensity score weighting, to robustly account for the hundreds of heterogenous categories and features exhibited by artworks, and thus provide a methodological advance for studying inequality in complex product markets. I also effectively control for structural factors that have driven much of previous work on markets. As such, I provide a uniquely complete picture of socioeconomic life in the contemporary art market. One that accounts for the relationship between products, producers, and market structure.

Third, I replicate and expand on the findings of emerging literature, which suggests a role for social status characteristics in the valuation of product markets (Tak et al., 2019). Importantly, I do so with real world market data, as opposed to the synthetic data that has previously been relied on. And finally, while some of sociology's seminal work on markets have involved the study of art (Padgett & Ansell, 1993; H. C. White & White, 1993), analysis terminates around 1945. Not only have there been

dramatic shifts in the structure, content, and practices of the art market since 1945 (Zolberg, 2015), but the exclusion of contemporary art has prevented the opportunity to address inequality among creative professionals working today, over 4 million of whom work in the United States alone (National Endowment for the Arts, 2019).

## 2. Case: the contemporary art market

In this manuscript, I use the case of gender inequality in the contemporary art market to study the influence of producer characteristics on market production and valuation. This case is ideal for several reasons. First, the market for contemporary art—i.e. art produced today, or in recent decades—is characterized by a diverse, complex, and innovative set of objects that generally lack a standard rubric for assessment and are exchanged without regulation (often privately) between buyers, sellers, and intermediaries who struggle to gain complete or symmetric information (Karpik, 2010; Velthuis, 2003, 2005). This results in a high degree of uncertainty toward the quality, and therefore the value, of goods (Espeland & Sauder, 2007; Fine, 1996; Greenfeld, 1988; Lang & Lang, 1988; Podolny, 1994; Posner, 1993; Zolberg, 2015). Without such certainty, standard economic theory would suggest that actors either refrain from participating in a market (Beckert, 2009; Koçak, 2003), or turn to other signals to justify value (Akerlof, 1970; Podolny, 1994; Spence, 1974; Stuart et al., 1999a). Yet while such non-market signals can unify collective choices and stabilize markets (Spence, 1974), they have also been found to increase inequality between products and diminish the predictability of product success (Lynn et al., 2009; Salganik et al., 2006). Here, I explore gender as one such “non-market” characteristic.

The contemporary art market is also an ideal case to study the influence of producer characteristics on product valuation and market structure because the

products of this market are often deeply connected with the identities of their producers. This is true of other products, such as entrepreneurship (Thornton, 1999), scholarship (Simcoe & Waguespack, 2010), and restaurants (Gaytán, 2008), but is acutely pronounced in art (DiMaggio, 1987). Defined as both a peculiar and a sacred good, the extreme view of art is that it cannot be translated into monetary value without degrading its essence (Zelizer, 2000). In this sense, art is not only valued by metrics of worth beyond money, but actually should not be valued by money at all (Boltanski & Thévenot, 2006). Whether we should, however, we clearly do, and for centuries have, monetized art. In the past twenty years, the art market has grown by over 200%, reaching a high of \$67 billion in 2018 with an annual turnover of approximately \$2 billion (*The Art Market 2019*, 2019). So, like other peculiar and sacred goods—nature (Fourcade, 2011), children, and life itself (Zelizer, 1979, 1981)—the monetization of art has developed into a highly rationalized market guided by specialized intermediaries, norms, and technologies (Velthuis, 2005).

Finally, the art market provides a useful case for the study of producer characteristics in product markets because of a present and long-standing history of gender-based inequality. Women have always been artists, yet historically they have also been barred from training, excluded from the canon, and overshadowed by their male counterparts. Indeed, no solo female artist was mentioned in the flagship art history textbook of United States until 1987 (Cowen, 1996). And today, despite popular

conceptions of the contemporary art world as a bastion of progressive practice, gender inequality remains pervasive and generally well-known among the insider art community. Women now comprise the majority of graduates with fine arts degree in the United States, and women are the growing majority of citizens who engage with fine arts as a form of cultural capital (DiMaggio & Mukhtar, 2018), yet social scientists have only just begun address the persistent inequality of this profession and industry. In her recent review of the sociology of art, for instance, Zolberg (2015) identifies four themes: the study of how institutions shape the emergence of art, investigations of artistic practice and consumption, demographics of public access, and epistemological inquiries into nature of “art” itself. The study of inequality is not included.

Despite this, there does exist a very small literature on gender inequality in art markets. It can be summarized as follows. Cowen (1996), in a historical analysis, articulates possible reasons for the lack of success among female artists and finds greatest support for discrimination against women and the effects of parenthood, which are biased against mothers. This work is not accompanied by statistical analysis. Cameron et al (2017), studies a sample of the prestigious Yale School of Art, and find that female graduates have significantly fewer auction sales but, when auctioned, obtain higher average prices. Similarly, Adams et al (2017), in an expansive study of auction markets, conducts two analyses, both of which support theories of discrimination against female artists. The first surveys over 1 million historical art auction sales and



finds not only that women's work sells for less (an average discount of 47%), but that this discount is exacerbated in nations with less gender egalitarian cultural norms. The authors then use experimental methods and simulated art to validate the observed devaluation of art by perceived gender of the artist. Finally, in their analysis of 2.6 million auction sales between 2000 and 2017, Bocart et al. (2018) find that female artists are much less likely to transition from the primary market, where art is sold for the first time through galleries, to the secondary market, where it is auctioned. The authors find that 96.1% of auction sales are for artwork by men. When women do have work auctioned, the authors find a 10% discount for their work, which rises to a 20% at higher value brackets of the market. At the very top of the auction market, where 40% of overall sales value is concentrated, there are no female artists.

The findings of these recent studies are informative and foundational for the present research, but they are also bounded in several respects. Most importantly, all prior research has evaluated auction sales, but only a small fraction of living artists—most of whom are relatively established in their careers and an especially small fraction of whom are women—participate this secondary market. By contrast, a much larger proportion of artists at most career stages sell in the primary market. The primary market now constitutes upwards of 62.5 of art sales worldwide (Pownall, 2017). The present study is thus unique in its access to and analysis of primary price data for contemporary artworks. It represents the first scientific study of valuation in the primary

market.

With the exception of Bocart et al (2018), who relied on 276 unspecified artistic characteristic to describe artworks in their analyses, prior studies have also lacked detailed information on artworks, and on other important information about artists, such as their level of fame and gallery representation. As described in further detail in the next section of this manuscript, the present research relies on data from 255,887 contemporary artworks by 18,624 artists, each of which is coded by art historians using a dynamic set of over 1,000 aesthetic, conceptual, and material features. This unprecedented data is drawn from the online art platform Artsy.net, whose Art Genome Project entails a taxonomy uniquely suited to describe contemporary trends of art.

### **3. Study A**

The first study of this manuscript addresses the following questions: do women make art with different characteristics than men? And, is art made by women valued less than art made by men? Not only are these foundational questions for the study of gender inequality in the art market, but they also provide an example for the broader investigation of how the social status characteristics of producers can influence product markets and their outcomes. I begin by discussing the concept of product quality and the need to consider the categories and features of products brought to markets when investigating inequalities between producers. I then distinguish the notion of status as generally used in the literature on product markets from the notion of status as generally used in the literature on labor markets. I argue that the latter is a means to formulate hypotheses about the salience of gender in contemporary art market. I then introduce my data and methods, test hypotheses, and evaluate results.

#### ***3.1 Product markets, categories, and features***

The question of how products in a market are valued has been, unsurprisingly, central to the research of economist and economic sociologists. Many have pointed attention to the notion of quality. Defined broadly, quality is the attributes of products and labor that are considered desirable in a market (Gould 2002, p1153; Olliver 2004, p. 148). Individuals can hold personal beliefs about quality, but the concept is generally understood as a more or less shared assessment and classification. It is arrived at

intersubjectively, via comparison, and is provisioned both by past perspectives and present cognitive frames (Beckert, 2019; Benjamin & Podolny, 1999). Hence, the notion of quality is contingent. Or such is the claim held as far back as David Hume (Hume, 1757). Often, however, neoclassical economic studies have simplified arguments of supply and demand by assuming that quality, as well as the consumer preferences that derive from it, are given and relatively stable (M. D. White, 2017). Behavioral economics has worked to insert a more realistic notion of human psychology (Frerichs, 2019), but the field of sociology has done the most to challenge this simplification. At scale, the discipline of sociology has combined the study of human cognition with understandings of social embeddedness and power to provide a holistic view of economic life.

For example, much of sociological inquiry into markets has attended to the social processes, institutions, and conventions that give rise to shared ideas of what is valuable (i.e. to quality) (Beckert, 1996). In so doing, scholars have often turned to the process of classification at the root of value judgments (Fourcade & Healy, 2017; M. C. Pachucki, 2012). This focus has, in turn, burgeoned a literature on the attributes of products themselves. Repeatedly, studies have demonstrated that market actors organize, interpret, and identify with products according to their relative attributes (H. S. Becker, 1982), including both the categories to which products belong (Caves, 2000; DiMaggio, 1987), and the underlying features of those categories (Tversky, 1977). This distinction is articulated well by Askin and Mauskapf (2017) and summarized here.

Categories in a market are largely agreed-upon conventions associated with groups of products that help people to organize and understand those products (Hsu, 2006; Jones et al., 2011). In cultural markets, such as those for film (Hsu, 2006), painting (Wijnberg & Gemser, 2000), food (Kovács & Hannan, 2015), literature (Frow, 1995, 2014), architecture (Jones et al., 2011), and music (Frith, 1996; Holt, 2007; Lena & Peterson, 2008), categories (e.g. genres) provide an order for producers and consumers as they shape identities and develop tastes (Bourdieu, 1993; Bryson, 1996; DiMaggio, 1987, 1997; Lizardo & Skiles, 2016; Peterson & Anand, 2004). They also supply a general basis for evaluation and quality assessment, as determining what an product is constitutes the first step in determining what it is worth (Lamont, 2012).

Yet useful as categories are, they yield only a general, even heuristic view of products and markets. Moreover, they have a tendency to blur, shift, and develop a significant degree of multivalence over time and throughout communities (Holt, 2007; Lena, 2012; Sonnett, 2004). By contrast, features of products are the underlying qualities that detail their composition. For example, a song may be categorized in the Country genre, but exhibit features of trap rap, twang, and conspicuous consumption; likewise, a painting may be categorized as Abstract, but entail features that are scientific, curvilinear, pastel, and mystic.<sup>1</sup>

---

<sup>1</sup> Hilma af Klint was a Swedish artist, now recognized as producing the earliest examples of Abstract painting. Convinced the world was not yet ready to understand them, Klint specified that her artworks should be kept secret for at least 20 years after her death. In 2019, an exhibit of her work at the Guggenheim

In some instances, a consistent set of features may aggregate into a category or genre, but that concordance is irregular, far from reliable and at times contentious (Anderson, 1991; Pontikes & Hannan, 2014).<sup>2</sup> With the onset of postmodern mentalities in many contemporary arts, the breakdown of categorical boundaries has become intentional. Likewise, the somewhat recent arrival of online platforms for cultural consumption—including Spotify for music, Netflix for film, and Artsy for art—has opened new opportunities for producers and consumers to (re)organize products in myriad and sundry ways, unconstrained by the physical and centralized spaces of yesteryear consumption (Petrusich, 2021). These new realities exemplify the theory that features of products provide an alternative means to categories for structuring markets. Both producers (Bourdieu, 1983) and products (Askin & Mauskapf, 2017; Kovács & Hannan, 2015) can be grouped in a conceptual space, according the features they do (and do not) have in common (Cerulo, 1988; Jones et al., 2011; Lena, 2006; Nunes & Ordanini, 2014; Peterson, 1992; Peterson & Kern, 1996; Tversky, 1977).

In the past, both categories and features of product have been relied upon in research to conceptualize the structure of markets and explain how consumers build

---

in New York City became the most attended exhibition the museum's history, drawing more than 600,000 visitors.

<sup>2</sup> Lil' Nas X's 2019 crossover song "Old Town Road", which exhibits features of American heartland-themed lyrics, trap rap, rock guitar styling, and hip-hop remixing, spent 19 weeks atop the US Billboard Top 100 chart—the longest-running number-one song in the chart's history to date. It also raised controversy when Billboard removed the song from the Country genre, sparking broadscale debate about gatekeeping and racism in the Country genre.

notions of quality and preference (Bourdieu, 1984; Bryson, 1996; DiMaggio, 1997; Goldberg, 2011; Halle, 1996; Lizardo & Skiles, 2016). They have also been used to predict the success of products and producing organizations according to their categorical niche (Hsu et al., 2009; Kovács & Hannan, 2015; H. C. White, 1981) or position in conceptual space (Askin & Mauskapf, 2017; Godart & Galunic, 2019). Yet what these studies have not done, or done very little, is seek to understand the role of categories and features in the inequalities that exist between market producers on the basis of social demographic characteristics, like gender.

Prominent theories of cultural consumption purport an identity alignment between the social categories of consumers and the categories (and features) of goods they choose to consume. At their origin, many of these same theories apply similar reasoning to an alignment between producers and what they choose to produce. From Marx to today, cultural theory and empirical analysis has insisted that the production of cultural elements in a society derive from the social systems in which they emerge and are distributed (Durkheim, 1915; Lena, 2006, 2012; Marx, 1867; Peterson & Anand, 2004). As such, when scholars investigate market production patterns, they can theoretically expect to find a specifiable relationship between the social categories ascribed to groups of producers and the categories ascribed to the goods they produce—especially symbolic goods (Bourdieu, 1993). This expectation can lead to hypotheses regarding the relationship between producers and their products. How, for example, do the social

characteristics of producers relate to the qualities of products they choose to produce? And, to what degree do the social characteristics of producers moderate the success of their products in a market, independent of the qualities they choose to produce? While these appear as straightforward questions, most market research has sidestepped the individual characteristics of producers, including identity markers of status, like gender, that are so often the bread and butter of sociology. Instead, research has investigated market inequalities via a form of status that derives from events taking place in the market. Most often, this conceptualization of status is measured at the level of organizations, not individuals. I will discuss the strengths and shortcomings of this approach to market inequality now.

### ***3.2 Product markets and status***

Compared the notion of quality, which is (again) the attributes of products and labor that are considered desirable in a market, the notion of status is a positional element of social structure (Washington & Zajac, 2005). In his foundational conceptualization, Max Weber (1978, p. 932) defined status as a social position based on honor or esteem and rewarded by privileges. Status often coincides with economic power and merit-based awards, but it is theoretically distinct. Although not a strict rule, the concept quality is generally operationalized as a rating in scientific literature and applied to products or labor, while status is generally operationalized as a ranking and applied to individuals, organizations, or social roles.



In market research, the notions of status and quality are often closely coupled. As status is generally more public and requires less domain knowledge to interpret, it is often theorized to act as a signal for the underlying quality of a producer's goods (G. S. Becker, 1964; P. M. Blau, 1964; Homans, 1961; Lynn et al., 2009; Spence, 1974). Indeed, status signals are especially relied upon when direct assessments of product quality are difficult to obtain; for example, when products are diverse, complex, or innovative, when a standard rubric for assessment is lacking, or when exchange is carried out among actors who lack complete and symmetric information (Beckert & Rössel, 2013; Espeland & Sauder, 2007; Fine, 1996; Greenfeld, 1988; Lang & Lang, 1988; Podolny, 1994; Posner, 1993; Zolberg, 2015). In such circumstances, uncertainty about the quality, and therefore the value, of goods is high. This uncertainty encourages actors either to refrain from participating in the market, or to look beyond products for the legitimization of value—including, commonly, to the status of producers (Akerlof, 1970; Podolny, 1994; Spence, 1974; Stuart et al., 1999a).<sup>3</sup>

Of course, if status is to be used as a signal for quality, then the two must be tightly correlated. Given that status itself is simply relative social prominence, its decoupling from quality can result in a market where products are not valued for their quality, but only by the social position of producers (Aspers, 2008). Phrased by Lynn,

---

<sup>3</sup> Even in the absence of quality uncertainty, scholars have found that concern with third party inference (what others think) can place more emphasis on status and obscure the market relationship between quality and value (Correll et al., 2017).

Podolny and Tao (2009, p. 756), *"[i]f each actor's attained status is a strong reflection of his or her quality, then we conclude that there is comparatively little social construction in the [market]. Conversely, if attained status is a weak reflection of quality, we can conclude the opposite."*

Notably, this quote neglects the potential for quality itself to be a social construction, but it nevertheless emphasizes the distinction between market judgments of value based on product attributes (quality) and those based on producer attributes (status).

The benefits of having high status in a market have mostly been analyzed at the level of organizations. Findings suggest that high status organizations are rewarded with relative ease gaining starting capital (Benjamin & Podolny, 1999; Jensen, 2008; Stuart & Ding, 2006), lower costs and barriers to entry (Phillips, 2001; Podolny, 1993a), a likelihood of higher revenue for a given level of quality (Benjamin & Podolny, 1999; Fombrun & Shanley, 1990; Kim & King, 2014), and an overall increase in survival (Baum & Oliver, 1991; Park & Podolny, 2000; Phillips, 2001; Podolny & Phillips, 1996; See also Sauder et al., 2012). These benefits, in time, can create a Matilda (1988) or Matthew (1988) effect, in which past privileges propagate onto to present means and judgments, thereby widening the gap between those who have and those who have not.

Measurement of status, in the context of product markets, has been strongly influenced by the work of Podolny (1993a). He outlines a model of market competition in which producer status is operationalized on the basis of exchange relations. This includes ties to arbiters and critics, affiliations with other producers, and connections to

important consumers (P. M. Blau, 1964; Elias & Scotson, 1965; Podolny, 1993b; Podolny & Phillips, 1996; Sauder, 2006; Stuart et al., 1999b; Torlò & Lomi, 2017; Washington & Zajac, 2005). In this way, status is operationalized on the basis of events occurring within the market, with the assumption being that birds of a feather flock together, and so high status producers will associate with high status alters.

While this approach to measuring status has provided many important insights into the market dynamics, it captures only market-derived elements of the social position originally theorized by Weber and others in the notion of status. It thus has the tendency to ignore the possibility that characteristics ascribed to producers prior to or independent of market events may also play a role in their status rewards (or penalties). Indeed, studies of “non-market” producer status characteristics, like gender, have been notably rare in product market literature, even as the incredibly few studies to address them have demonstrated significant variance in market outcomes by these characteristics (e.g. Kricheli-Katz et al., 2019; Kricheli-Katz & Regev, 2016; Tak et al., 2019). Interestingly, the study of such characteristics is more prevalent in the adjacent literature on labor markets, where the particular subfield dedicated to explaining the gender gap in wages has provided prolific insight.

### ***3.3 Labor markets and status***

Historically, scholars of the labor market have explained the gender wage gap along two lines of reasoning: supply and demand (Blackburn et al., 2002; F. D. Blau &

Kahn, 2007; England, 2005). Supply-side perspectives—often associated with the research of economists—focus on the behavioral patterns of women as potential labor force participants. Stated crudely, these perspectives argue that women earn less than men because they select into different occupations. They address issues of market segregation. By contrast, demand-side perspectives emphasize biased and discriminatory patterns of behavior among employers, colleagues, and consumers. They address issues of comparable worth. Stances from this latter perspective are often taken by sociologists, although there is a long lineage of scholarship across disciplines focused on gender, as well as race (Notably, G. S. Becker, 1957; Du Bois, 1899).

Although supply and demand theories of inequality in the labor market have at times been positioned as conflicting, they are not generally envisioned as such by the scholars investigating them. Those employing demand-side perspectives are quick to acknowledge the reality and the complexity of supply-side arguments. While generally resisting human capital and rational choice explanations of occupational segregation (which emphasize women’s actual or expected motherhood responsibilities as a leading cause of their job selection), demand-side scholars recognize that the general argument that women and men select into different occupations is robust (Blackburn et al., 2002; F. D. Blau & Kahn, 2007; England, 2005). Indeed, taken together, several segmentation explanations posited in the supply-side literature account for upwards of fifty percent of the gender wage gap in the United States (F. D. Blau & Kahn, 2016). Even after

controlling for such things as industry, part-time work status, work experience, seniority, training, and other purportedly supply-side variables, however, a gap in the wages earned by men and women remains (Budig & England, 2001; Dinovitzer et al., 2009; Lundberg & Rose, 2000; Waldfogel, 1997; Weichselbaumer & Winter-Ebmer, 2005). It is at this juncture that demand-side mechanisms are brought to bear (Reskin, 1993).

Advocates of demand-side perspectives argue that discriminatory practices and gender-based norms account for a non-trivial amount of the gender wage gap left unexplained by supply-side arguments, and in some cases may even be their origin (Bergmann, 1986; Correll, 2004; Reskin & Roos, 1990). Drawing on a foundation of expectation states theory (Berger et al., 1972)<sup>4</sup>, demand-side studies often address gender as a diffuse social status characteristic, meaning that it is a nominal characteristic with widely shared belief about the competence, ability, or esteem of categorial members (Berger et al., 1972; Correll et al., 2017; Correll & Ridgeway, 2006). These beliefs have been found to influence expectations and behaviors even when the characteristic upon which they are grounded is not directly relevant to the task or context at hand (Berger & Fişek, 2006; Ridgeway, 2001).

Unlike the notion of market status previously discussed, which primarily derives from events and associations taking place inside the market, the concept of a diffuse

---

<sup>4</sup> Status Characteristics (Ridgeway & Correll, 2004) and Socially Endogenous Inference (Correll et al., 2017; Gould, 2002; Lynn et al., 2009; Podolny, 1993a; Salganik et al., 2006).

social status characteristic derives from the broader social context in which markets are embedded. These status characteristics are ascribed qualities of self that producers bring to the market, and for which they face bias, privilege, or discrimination. In the case of gender, status beliefs often downwardly bias market judgments of women's work, resulting in systematically different outcomes for otherwise similar women and men (Ridgeway & Correll, 2004; Ridgeway & Smith-Lovin, 1999; Wagner & Berger, 1997).

Field experiments, including audit and correspondence studies, have provided evidence for demand-side arguments of the gender wage gap. These studies demonstrate that employers in many fields discriminate against women in job applications (Riach & Rich, 2002), and further indicate that women face different expectations and judgments than men by managers, co-workers, and consumers with regard to their competence, productivity, ability to innovate, and leadership style (Brescoll & Uhlmann, 2008; Eagly & Karau, 2002; Foschi, 2000; Goldin & Rouse, 2000; Gorman & Kmec, 2009; Heilman, 2012; Moss-Racusin et al., 2012; Proudfoot et al., 2015; Steinpreis et al., 1999). This is especially true in situations that are traditionally male-dominant or ambiguous in standards of judgment (Correll et al., 2017; Ko et al., 2015; Steele et al., 2002).

As a diffuse status characteristic, gender is primed in most social contexts, but the strength of its influence is not fixed. Gender combines with other identities and roles in situ, with the intersection determining the ultimate behavior of individuals and their

alters (Ridgeway & Correll, 2004; Wagner & Berger, 1997). At times, gender may be hardly noticeable or overshadowed by additional social characteristics; at others, it may be accentuated or entirely the point (Ridgeway & Smith-Lovin, 1999). For this reason, scholars refer to gender as “effectively salient” when gender stereotypes are culturally linked to activities of a given context and influential upon their expectations, judgements, and outcomes (Correll, 2001). If there exist systematic differences between the value of work done by women and men in a market, for example—net of other relevant qualities upon which value may be assessed—then a strong argument can be made that gender has become an effectively salient status characteristic to the activities of that market.

### **3.4 Hypotheses**

In what follows, I formulate two hypotheses that expand on the literature previously discussed to investigate the nature of gender disparity in the market for contemporary art. First, I consider the relevance of product categories and features to the organization and gender differentiated outcomes of the art market. This leads to a supply-side hypothesis that is also a colloquial intuition about the gender disparity of the art world: that women and men choose to produce art with different characteristics. For this argument to be upheld, the following hypothesis should be supported.

*Hypothesis 1: Women and men make artwork with different characteristics.*

An alternative explanation, reflecting a demand-side perspective, is that gender is an effectively salient status characteristic of the art market. Women, by this argument, face downwardly biased judgments even when producing art with the same categories and features as men. For this argument to be upheld, the following hypothesis should be supported.

*Hypothesis 2: Net of artistic characteristics, artwork by women is valued less than artwork by men.*

As described in further detail below, I propose to test these hypotheses using a new dataset of contemporary artworks that have been linked to artists and coded in detail for the aesthetic, conceptual, and material features they depict.

### ***3.5 Part 1: Do women make art with different characteristics than men?***

To assess my first hypothesis, that women make art with different characteristics than men, I implement machine learning classification and test how well an algorithm, provided with the salient qualities of an artwork, can be trained to correctly attribute that work to either a female or male artist. I also descriptively explore the patterns of aesthetic, conceptual, and material features used by female and male artists. Together, these analyses provide the best estimate to date of whether women and men make art with different characteristics.



### 3.5.1 Data

In this study, I leverage data from Artsy (artsy.net), an online platform designed to facilitate the discovery and collection of art. As a company, Artsy has created a global database of over 1.8 million artworks, produced by over 130,000 artists. A portion of this database has been coded by the Art Genome Project, Artsy's proprietary classification system and the technological framework for connecting artists and artworks on the site. The Art Genome Project is driven by a dynamic taxonomy of over 1,000 art-relevant features (aka "genes"), which collectively describe artworks according to their disciplines, materials, physical attributes, styles and periods, object types, and geographic settings. These features are assigned to artworks by a team of engineers and art historians who draw on scholarship and artificial intelligence to determine values. Each feature is assigned along with a score from 0 to 100, thus capturing how strongly the feature is represented in the work.<sup>5</sup>

Although the Art Genome Project has coverage from the prehistoric to the present and a global scope, it is especially designed to accommodate the rapid evolution

---

<sup>5</sup> Interestingly, in the 1960s, philosopher Artur Danto (1964) argued that the history of art could be envisioned as a collection of aesthetic possibilities that artists have chosen to depict through the features of their work (Rule & Bearman, 2015). If every one of the possibilities were to be listed along the columns of a grid, then a given artwork could be defined by iterating through each possibility and marking a 1 under every quality that the work entails, and a 0 for every quality that it does not. Completing this task for every artwork would result in a matrix of artworks (rows) by features (columns); with each new artwork a new row would be added, and with each new artistic possibility a new column would be added. As a whole, this matrix would define art. The data used in the present study mimics that structure.

of the contemporary art world and includes such features as: African Diaspora, Blurred, Censorship, Digital Art, Eye Contact, Fluxus, Grotesque, Humor, Impasto, Japonisme, Kitsch, Linear Forms, Mosaic, Nude, Op Art, Personal Histories, Related to Music, Self Portrait, Taxidermy, Urbanization, Violence, and Watercolor. This provides the Art Genome Project with the unique benefit of encompassing both categories and features of art, as described earlier in this manuscript. For my analyses, I subset the Artsy data to include only contemporary artworks by solo artists, and to exclude artworks not coded by the Art Genome Project team of art historians and engineers. I also exclude artworks that are primarily comprised of video, digital, or performance elements, as these forms of art are difficult to compare and monetize.<sup>6</sup>

The result is a dataset of 255,887 contemporary artworks by 18,624 artists represented by 2,284 galleries. For the majority of artists in this sample (71%), gender information is provided by the gallery listing the artwork on the Artsy site. Missing gender information is imputed using a combination of hand-coding and census record distributions for first name (see Appendix B for more detail). I also conduct sensitivity analysis using list-wise deletion of artworks without gallery-provided gender data. Results are not substantively altered by this adjustment.

---

<sup>6</sup> Further detail on the Art Genome Project, inclusion criteria can be found in Appendix A.

### 3.5.2 Methods

Machine learning classifiers are programs that predict which category of a dependent variable individual cases in a dataset belong, based on their values for other features. The use of classifiers has swelled with both the advancement of computational power, and the size and complexity of available datasets. They are now implemented in countless areas of modern life, and often applied in scenarios with large, complex datasets, such as genomics, computer vision, drug discovery, social media, and speech recognition. Social scientists have been slower to adopt machine learning methods. In large measure, this is due to concern over the “black box” nature of their processes, in which high prediction rates can be obtained without the ability to explain how (Fourcade & Healy, 2017). Nevertheless, in recent years, several prominent calls for the integration of machine learning with more traditional hypothesis-testing approaches have been made (Mason et al., 2014; Mullainathan & Spiess, 2017; Yarkoni & Westfall, 2017). The present study provides one example of this.

There are many types of machine learning classifiers, each of which uses a different algorithm to efficiently search possible parameter values until correct predictions are optimized (i.e. until the loss function is minimized), net of any regularization. For the present study, I implement random forest classification (Breiman, 2001). Random forest is a type of decision tree classification, meaning that the algorithm builds a set of iterative decision rules that place individual cases into categories of the

outcome. The cases in the present study are artworks, with decision rules based on the Art Genome Project features previously discussed. The categories into which artworks are classified are ‘female’ and ‘male,’ indicating that the artist of the work is female or male.

Random forest classifiers are one of the most general-purpose algorithms in modern machine learning because they are non-parametric, fast to train, hard to over-train, and robust to outliers and noise (Biau & Scornet, 2016; Folleco et al., 2008; Touw et al., 2013). They also generally produce high prediction accuracy and account for interactions among variables without the need to specify them directly. Like most classifiers, random forests are a form of supervised learning, meaning researchers split data into two sets—training and test—prior to analysis. The training set is used to train the classifying model, and the test set is used to test how well the model predicts the outcome for data it has not yet observed. I train a classifier on artworks by 70 percent of artists in the dataset ( $n = 183,111$  artworks) and test on artworks by the remaining 30 percent of artists ( $n = 76,489$  artworks).<sup>7</sup> I implement 10-fold cross validation and down-sample to account for modest imbalance in the outcome.

---

<sup>7</sup> The reason for sampling on artists rather than artworks relates to the nature of random forest classifiers and the aim of the present research. Were the classifier to be trained on works by artists whose work is also in the test set, it is likely that correlation of artworks by the same artist would drive prediction. The classifier’s accuracy would then answer the question, “can a classifier learn to identify the repertoire of an artist based on past works and impute that artist’s gender onto new works?” This is not my research question. As previously stated, I aim to test whether—without prior information about an artist—a classifier can accurately predict the artist’s gender based on the qualities of their work. For that reason, I train on works produced by a different set of artists than I test on.

Given that classifiers are often applied to very complex datasets with many variables and a large number of cases, it is common to perform variable (i.e. feature) modification and selection prior to training. In cases where the goal is to optimize the accuracy and efficiency of the classifier (for instance, when the classifier will be used regularly for online tools or business) this process includes removing, combining, and coarsening predictive variables so that redundant and overly intricate information is reduced. Doing so allows the classifier to run quickly while still performing accurately.

For the present study, such computational efficiency is not so important. As such, the only feature modification taken was to binarize values of the Art Genome Project applied features, and the only feature selection was based on theoretical and statistical, rather than computational, needs (see Appendix B for more detail). This variable selection and modification results in 594 art-relevant features, all of which are used in the primary random forest classification to predict artist gender. Due to the proprietary nature of the Artsy dataset, a detailed list of these feature is only available upon request to the author. As of June 2021, general information can be found at [artsy.net/categories](https://artsy.net/categories).

As seen in Tables 1 and 2 below, the median number of artworks that a feature is applied to in this sample is 2,115. The median number of features applied to an artwork is 12, for both female and male artists. The median year of production is 2013, demonstrating significant skew toward more recent work. Artists have a median of 8

works in the dataset, with a minimum of 1 and a maximum of 1,581. Female artists comprise approximately one third of the artists in this dataset (female n = 6,328; male n = 12,296) and produce approximately one-third of the artworks (female n = 77,153; male n = 178,734). Such proportions mirror those reported in generalized surveys of the art market (*The Art Market 2019*, 2019).

**Table 1: Number of artworks ascribed to each feature**

	All artworks	Artworks by Women	Artworks by Men
<i>Min</i>	54.0	5.0	36.0
<i>Median</i>	2,115.0	594.0	1,441.5
<i>Mean</i>	5,328.0	1,634.7	3,693.7
<i>Max</i>	115,310.0	35,320.0	79,990.0

Artwork n = 255,887; Artist n = 18,624; Gallery n = 3,384

**Table 2: Number of features ascribed to each artwork**

	All artworks	Artworks by Women	Artworks by Men
<i>Min</i>	5.0	5.0	5.0
<i>Median</i>	12.0	12.0	12.0
<i>Mean</i>	12.4	12.5	12.3
<i>Max</i>	45.0	45.0	37.0

Artwork n = 255,887; Artist n = 18,624; Gallery n = 3,384

Before moving to the results of classification, a brief note on what precisely I aim to measure is appropriate. As pointed out by Mullainathan and Speiss (2017), machine learning methods provide new tools for different problems in the social sciences. Classifiers in particular are generally useful for prediction and ill-equipped for the type

of parameter estimation that may help to explain how a dataset was created, which is the common aim of inferential statistics. Among other things, however, classifiers can optimally determine whether patterns in data are strong enough to correctly place individual cases into specified categories. My empirical aim in using a classifier is thus not to mimic the human mind standing before a work of art, with all the context and additional information that may come with such an experience. Nor is it, in detail, to explain which features of art might distinguish between women and men. Rather, more simply, I aim to determine whether women and men make art with different characteristics.

### **3.5.3 Results**

#### **3.5.3.1 Artwork-level random forest classification**

The performance of the random classifier is a balanced accuracy of 0.62 (AUC-PR = 0.5), indicating that there is a 0.62 probability that the classifier will correctly classify an artwork by the gender of the artist, given the features of the work.<sup>8</sup> Balanced accuracy averages the proportion of correct predictions for both female and male artists, thus accounting for any imbalance in the overall gender proportions of the dataset (Brodersen et al., 2010). With approximately one third of artists and artworks being female and one third of artworks being produced by female artists in this sample, there

---

<sup>8</sup> The optimal threshold for classifying was determined by the Informedness metric.

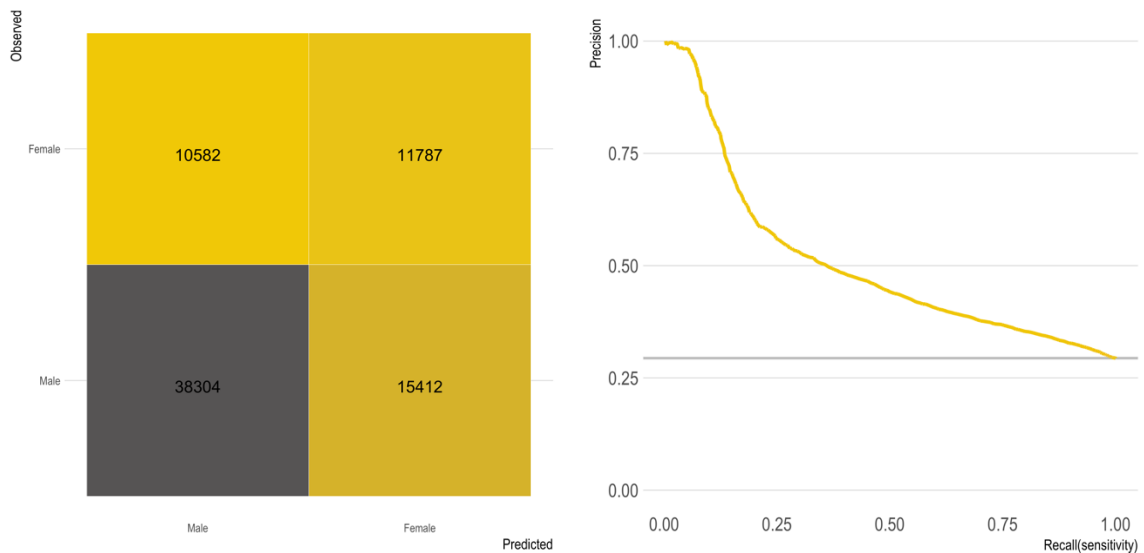
is relatively little imbalance to begin with and down-sampling also adjusted for this discrepancy. We can therefore be confident that results of classification are not merely a reflection of underlying gender distributions. A similarly unbiased metric, the phi coefficient, is positive but modest at  $\varphi = 0.23^9$ , as is Cohen's kappa (0.22), which indicates that the classifier performs only slightly better than the expected accuracy of classifier relying solely on the base frequencies of each group.

A decomposition of the overarching accuracy indicates that neither artworks by female artists (specificity = 0.71) nor artworks by male artists (sensitivity = 0.52) obtain a particularly high level of prediction accuracy. Figure 1 depicts the confusion matrix and precision-recall plots for further information. It shows that a significant number of false positives (predicting male artists to be female) is introduced in order to obtain perfect recall (identification of all female artists). Furthermore, when approximately 50% of artists predicted to be female are actually female, only approximately 25% of all female artists are identified.

---

<sup>9</sup> The phi coefficient (aka Matthew's correlation coefficient) takes into account true and false positives for both classes, as well as true and false negatives. As such, a high value (close to 1) indicates that both classes are predicted well by the classifier, even if one class is disproportionately under- (or over-) represented. By contrast, a low value (close to -1) indicates that both classes are predicted poorly.






**Figure 1: Confusion matrix and precision-recall plot of results from random forest classification**

The Precision-Recall curve (above, right) depicts precision on the y-axis, which is the number of true positives divided by the sum of the true positives and false positives, and recall (or sensitivity) on the x-axis, which is the number of true positives divided by the sum of the true positives and the false negatives. Precision-Recall curves are particularly good at assessing the correct prediction of the minority class (female), and are thus more informative than a traditional ROC plots when evaluating binary classifiers on imbalanced datasets (Saito & Rehmsmeier, 2015).

With minor variation, these results are robust to subset analyses for different mediums of art (i.e. photographs, paintings, textiles, sculptures, drawings, and other), different specifications (i.e. using non-binarized values for the Art Genome Project data,

or dimension reduction to reduce noise), and to the use of alternative classification methods, such as Logistic regression, Ridge regression, and Boosted trees (see Figure 2 below and Appendix C.2 for descriptions of each model and subset; more detail available upon request).



**Figure 2: Robustness checks of random forest classification.**

### 3.5.3.2 Artist-level descriptive statistics

A descriptive investigation of the artistic features used by female and male artists provides a fuller understanding of the random forest classification results. Figure 3 depicts features by the binary log of the ratio of relative frequencies for female and male artists (i.e. the log ratio). It contains information for 18,624 artists' work, dated 1945 to 2018. A value close to 0 on the y-axis indicates near equality in use of that feature

between female and male artists; positive values indicate that the feature is more frequent among female artists than male, and negative values indicate that the feature is more frequent among male artists than female. The right y-right-axis provides simpler interpretation of the log ratio, describing how many times more frequent a feature is among female artists and vice versa. The x-axis indicates the number of applications of each feature to artworks in the dataset.

This figure demonstrates that there are few instances of substantial difference between female and male artists in their use of individual artistic features. On median, the absolute log ratio is 0.28, indicating that the median feature is used 28% more among one gender than the other. Ninety one percent of features ( $n = 553$ ) fall below an absolute log ratio value of 1, which is the threshold at which one gender is twice as likely as the other to use an artistic feature. There are, however, 51 features that have an absolute log ratio at or above 1. Feminist art (log ratio = 5.5), for example, is the most likely to be produced by a female artist. Other more female-typic features include knitted, crocheted, embroidered, and woven (between 1.3 and 3.0), manga art (1.3), and art engaging with the themes of gender (2.0), gender and sexual identity (1.6), fertility, pregnancy and birth (1.2), and domesticity (1.2). By contrast, engineering and

construction design (-2.3), graffiti and street art (-1.5), and graphic and poster design (-1.5 and -1.8) are features over-represented in artwork by men.<sup>10</sup>

Fifty-one percent of features with an absolute log ratio value above 1 occur more in the artwork of female artists, and of the top ten most extreme values, 50% are more male-typic and 50% are more female-typic. Notably, features used disproportionately by one gender or another are used by a relatively low number of artists, with a median of only 91 artists for features with an absolute log ratio of 1 or more, compared to a median of 433 artists otherwise. This may help to explain why classifying algorithms were unable to rely on these gender imbalanced features alone to distinguish between artworks by women and men.

---

<sup>10</sup> It is important to note that design objects—or, functional objects that are appreciated for their aesthetic qualities but are not art per se—are excluded from this analysis. As such, these statistics can only be interpreted as patterns of categories and features as they appear in artwork, not design objects. Graphic design and posters not defined as art, for example, may be less overwhelmingly produced by men.



### 3.5.4 Discussion

The combination of machine learning and descriptive statistics provided in this study indicate substantial overlap in the types of art created by women and men, but also significant difference in certain qualities. Given the unprecedented detail of the data analyzed, with 255,887 recent artworks described by over 1,000 art historian-coded categories and features, results provide the best estimate to date on whether women and men make art with different characteristics. The inability to rely on a robust, broadly implemented method of machine learning to distinguish between art produced by women and men withholds support for hypothesis 1, that women and men make art different characteristics. Thus, unlike studies of the labor market, where upwards of 50% of variance in women's and men's wages can be explained by selection into different occupations (F. D. Blau et al., 2013; F. D. Blau & Kahn, 2016), the findings of this analysis suggest that market segregation is unlikely to be a leading factor in gender gap for listing prices in the primary art market. The next section addresses price differences directly.

In the past, machine learning algorithms, such as those used in the present study, have been found to perpetuate and amplify human biases, including those related to gender and other social identity markers, like race (Bolukbasi et al., 2016; Buolamwini & Gebu, 2018; Kleinberg et al., 2018; Mullainathan & Obermeyer, 2017). This occurs when the data on which algorithms are trained hold such bias and no effort is made by

researchers or engineers to generate informed benchmarks and adjustments. Indeed, algorithms applied naively or misinterpreted in their outcomes can—as with all statistical methods—lead to errant and potentially harmful repercussions. The relatively low accuracy of the classifier implemented in this study provides little indication that the human coders of the Artsy Art Genome Project produced data that was biased by gender in their application of artistic features. This should not be interpreted as an audit of such coder bias, however, as results speak only to overall patterns in the data, and not to specific features. More experimental work would be needed to investigate whether particular features were applied by human coders with bias, or whether actual differences between male and female artists were obscured.

Of course, the ability of algorithms to predict the gender of artists by the categories and features of their work is likely to vary across time periods and types of art. The present study results are robust when looking within principal genres of recent work, however—including paintings, photography, sculptures, drawings, and textiles. Furthermore, and despite broadscale evidence that the machine learning methods used in this analysis can distinguish complicated patterns in large and complex datasets, this is not the first study to find moderate levels of prediction when applied to social scientific questions (e.g. Salganik et al., 2019, 2020).<sup>1</sup> While there are substantive

---

<sup>1</sup> The Fragile Family challenge was an open call research initiative aimed at predicting six family outcomes using over 1,500 variables in a fifteen-year dataset collected by social scientists. The challenge received 160 academic submissions; many of which relied on machine learning methods, none of which made very

differences between the present and prior studies, the cases collectively leave scholars with an open question: is the lack of accuracy in prediction a reflection of reality, of methodological limitations (i.e. the patterns in the data exist but are not recoverable by the methods used), or of a limitation of the data (i.e. the data does not contain the information necessary to pattern the outcome). The conclusion of the present analysis thus remain moderate; it is that, in our best estimate to date, using the most detailed data and cutting-edge methods, women and men do not make art with systematically different characteristics.

While the random forest classification analyses of this study did not discover deterministic differences in the artwork of women and men, the descriptive statistics revealed that certain qualities of art are more common in the artwork of one gender of artist or the other. These distinguishing features are used by fewer of artists, compared to more gender balanced features, making their use as predictors at the population-level less useful. They are prevalent enough, however, to provide signal for further investigation into the interplay of gender, artistic categories and features, and market outcomes. In the part two of this study, and in Study B of this dissertation, I rely on

---

accurate predictions (Salganik et al., 2019, 2020). There are very substantive differences between the present research and research like the Fragile Families challenge. In particular, the aim of the latter was to recover patterns in data that lead to the patterns in the outcomes being predicted. The predictors precede the outcome. For example, how well does childhood family income relate to later levels of education? The present research, on the other hand, aimed to recover patterns in the data that may have been created by patterns in the outcomes being predicted. The predictors follow the outcome. For example, how well does the presence of curvilinear lines in art predict the gender of the artist? Yet while this difference is important theoretically, both cases engage the topic of how machine learning methods relate to social science.



variation in artistic categories and features to test hypothesis about the disparity of artwork listing pricing by the gender of artists and the aesthetic qualities of their work.

### ***3.6 Part 2: Is artwork by women valued less than artwork by men?***

Part 1 of this study indicated that women and men do not make art that is highly distinguishable by the aesthetic, conceptual, or material features they depict. The question of whether art produced by women is valued less than art produced by men—hypothesis 2—remains open. Previous research from the art auction market provides mixed evidence (Adams et al., 2017; Bocart et al., 2018; Cameron et al., 2017; Coate & Fry, 2012; Edwards, 2004), and an empiric assessment of primary sales data has yet to be undertaken. Below, I implement mixed-effect linear regression analysis to predict the listing price for contemporary artworks by the gender of artists, net of artistic characteristics and other contributing factors related to the value of art. I also implement quantile regression to explore differences in the influence of gender by price strata of the market.

#### **3.6.1 Data**

In these analyses, I again use the Artsy dataset, subset to artworks produced since 1999, and for which a gallery listing price is available. This sample comprises 102,908 artworks by 11,253 artists. The dependent variable, gallery listing price, is measured in US Dollars (USD). Artworks whose original price was listed in another currency were converted to USD using the average Central Bank conversion rate for

2017. As in the classifier analysis, artists' gender is measured categorically as female or male, with information on the imputation of missing data provided in the Appendix B. In addition to artists' gender, I control for a number of factors likely to influence the price of an artwork. These are of three types: those related to the artwork, those related to the artists, and a dummy for each gallery that listed artworks on the Artsy website.

Controls related to the artwork include the size of the work, measured in square or cubic inches, depending on if the work is two- or three-dimensional. Previous studies have found size to be a leading correlate for the price of art in auction sales (Velthuis, 2005). Missing data for the size of an artwork was handled via the listwise deletion of 2,522 artworks. I also control for the artwork's date of creation and, as discussed below, rely on propensity score matching to account for aesthetic, conceptual, and material qualities of the work.

Artist controls include three measures of visibility or fame (Menger, 1999). First, the number of followers on the Artsy platform; second, the number of galleries with which the artist has associated on the platform (i.e. the number of galleries that have listed at least one artwork by the artist); and third, the number of exhibition events listed for the artist on the platform between 2014 and 2018, with 2014 selected as the cut-off because it is the modal earliest event date in the dataset. As a measure of productivity, I also include the number of artworks listed by the artist on the Artsy platform.

Finally, fixed gallery-level characteristics are accounted for by the inclusion of a dummy for each of the 1,998 galleries listing artworks. As background, artists' careers typically begin with the completion of an art degree, after which many obtain gallery representation. A gallery, which is often a small business, acts as an agent with the goal of creating demand for the artist in exchange for a commission on the price of artworks sold. Gallery services may include provision of exhibition space, advertisement and journalism, support for logistics of art fair participation, relationships with collectors, and archival preservation. Depending on their success in these roles, galleries can provide considerable symbolic capital and opportunity for their artists. As such, the inclusion of a dummy for gallery in the proceeding analyses accounts for this structure of organization in the art world and any systematic gallery trends in the data.

### **3.6.2 Methods**

#### **3.6.2.1 Log linear mixed effect regression**

Given the multilevel nature of the data, with individual artworks (level 1) nested within artists (level 2), I estimate a log linear mixed effects regression that predicts the gallery listing price of artworks, with random intercepts for artists. I also include covariate balancing propensity scores, to encourage the robustness of comparison between artworks by women and men. These weights are described in further detail below. Random intercepts of all regression analyses capture the possibility that the average value of artworks cluster by artist, depending on artist-level attributes that are

unobserved. Thus, in hierarchical form, the equation for gallery listing price,  $P_{ij}$ , is as follows:

$$P_{ij} = \beta_0 + \beta_1 Z_j + X_{ij} + u_{oj} + e_{ij}$$

Where  $i$  and  $j$  index artworks and artists, respectively;  $Z_j$  denotes the vector of artist-level variables, including gender, and  $X_{ij}$  denotes the vector of artwork-level variables, in this case the size of the work and the year of production.  $u_{oj}$  is the residual variance component across artist intercepts and  $e_{ij}$  is the residual variance at the level of artworks. Due to significant skew, I take natural log of the outcome and of all continuous covariates.

### 3.6.2.2 Covariate balancing propensity score weights

In experimental studies, the act of randomization produces a state of statistical equivalence (i.e. “covariate balance”) whereby treatment and control groups are equivalent on both observed and unobserved covariates. By contrast, observational studies such as the present rely on data that has not been generated by randomization and may exhibit difference between treatment and control cases (i.e. women and men). Commonly, social scientists address such covariate imbalance by including possible confounders in regression models. In their influential article, however, Rosenbaum and Rubin (1983) showed that when treatment assignment is strongly ignorable given observed covariates, an unbiased estimate of the average treatment effect can be calculated by adjusting for the propensity score alone, where propensity score is the

conditional probability of receiving treatment, given covariates. Many methods have since been proposed for using propensity scores in observational analyses, and their use has become a fundamental part of research across the social sciences.

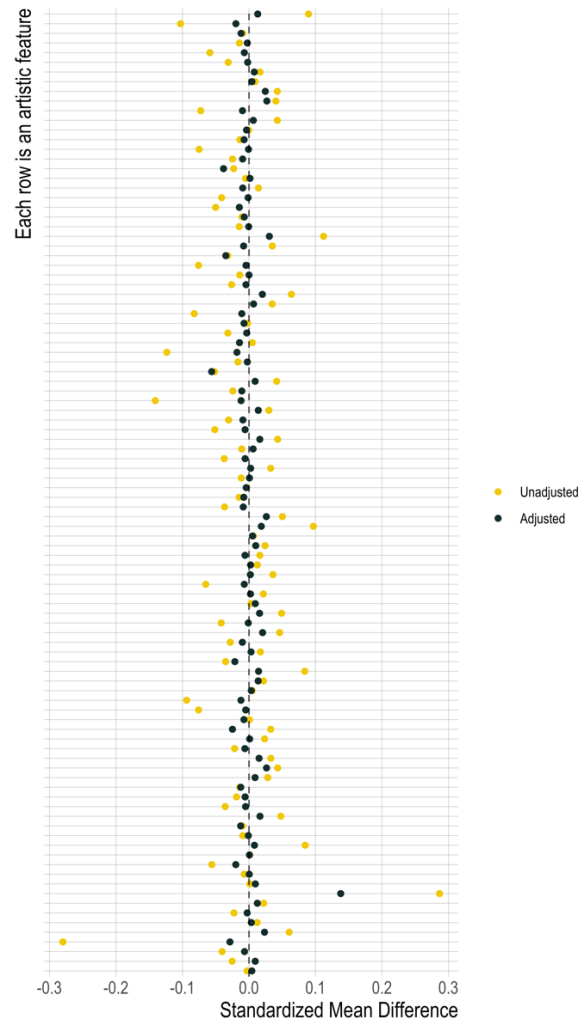
In the present analyses, I address the potential for imbalance in the use of certain artistic features between female and male artists by calculating average treatment effect on the treated (ATT) weights derived from the method of covariate balancing propensity scores (Imai & Ratkovic, 2014).<sup>2</sup> By regressing all of the aesthetic, conceptual, and material features of the Art Genome Project described previously on artists' gender, ATT weights adjust artworks by male artists to be more statistically comparable to artworks by female artists on these attributes. Specifically, artworks by female artists are given a weight of 1, because I am interested in the effect of being female on the listing price of artworks, while artworks by male artists are assigned weights that are greater than zero and vary as necessary to optimize balance. More weight is assigned to artworks by men

---

<sup>2</sup> Despite their prevalence, the practical difficulty of estimating propensity scores has presented scholars with a new problem: misspecification of the model used to estimate propensity scores using high-dimensional covariates can bias estimated treatment effects (King & Nielsen, 2019). Covariate balancing propensity scores is a proposed solution to this problem (Imai & Ratkovic, 2014). The method exploits the fact that propensity scores serve as both covariate balancing metrics and the conditional probability of treatment assignment. By combining a set of moment conditions with the standard estimation procedure for propensity scores, the method dramatically reduces the potential for bias in estimates and is more robust than misspecified logistic regression or boosted classification trees (Wyss et al., 2014). The method has been successfully used in recent studies related to attitudinal shifts toward gay and transgender rights (Flores et al., 2018; Flores & Barclay, 2016), the role of government finance in conservation metrics (Waldron et al., 2017), and the impact of remittances on poverty (Dey, 2015).

that have more comparable artistic features to artwork by women, and less weight is assigned to artworks by men with features that are more different.

Figure 3 provides descriptive statistics for the aesthetic, conceptual, and material features of art by female and male artists before and after balancing the samples with this method. It demonstrates that, prior to weighting, female and male artists are dissimilar in their use of several artistic features, on average. This underscores the conclusion of descriptive statistics in part one in this section, and affirms the need to adjust for these factors when estimating gender disparities in price. After weighting, there are significantly smaller differences between female and male artists. Due to space constraints, the figure plots a random sample (n=100) of the 594 features used in weight construction. The figure table reports statistics for all features.



	<b>Min</b>	<b>1<sup>st</sup> Q</b>	<b>Median</b>	<b>Mean</b>	<b>3<sup>rd</sup> Q</b>	<b>Max</b>
<i>Unadjusted</i>	-0.28	-0.03	-0.01	-0.01	0.02	0.80
<i>Adjusted</i>	-0.06	-0.01	0.00	0.01	0.01	0.24

**Figure 4: Plot of standardized mean difference in artistic features for artwork by women and men, adjusted and unadjusted by covariate balancing propensity score weights**

### 3.6.3 Results

#### 3.6.3.1 Descriptive statistics

Table 3 depicts descriptive statistics about the artworks and artists in this sample. The average listing price for artwork by women is significantly below the average listing price for artwork by men (\$9,814 versus \$16,619), though the gap in median prices is much smaller (\$4,000 versus \$5,500). This reflects a common gender-based trend in covariates, whereby the mean disparity between women and men presents a much larger gap than the median disparity, in every instance indicating that male artists have higher highs. In this way, the primary market is not substantially different than the auction market for contemporary art, despite the former including a broader set of artists and many more women (Bocart et al., 2018). Pairwise correlations of all variables are provided in Table 4.

**Table 3: Descriptive statistics for variables used in regression analysis**

		ARTWORKS		
		Overall	Female	Male
<b>Listing price (USD)</b>	<i>Mean</i>	14,226.5	9,813.5	16,618.9
	<i>Median</i>	5,000.0	4,000.0	5,500.0
	<i>Sd</i>	81,609.8	38,600.0	97,195.0
	<i>Min</i>	10.4	20.0	10.4
	<i>Max</i>	15,000,000.0	4,000,000.0	15,000,000.0
<b>Size (sq/cu in)</b>	<i>Mean</i>	6,086.3	5,136.3	6,601.3
	<i>Median</i>	900.0	746.9	988.9
	<i>Sd</i>	90,814.4	91,479.8	90,448.2
	<i>Min</i>	0.0	0.0	0.0
	<i>Max</i>	8,985,600.0	8,985,600.0	8,489,664.0
		ARTISTS		
		Overall	Female	Male



<b>Artist artwork count</b>	<i>Mean</i>	117.9	83.5	136.6
	<i>Median</i>	62.0	56.0	65.0
	<i>Sd</i>	270.3	102.9	325.5
	<i>Min</i>	1.0	1.0	1.0
	<i>Max</i>	7,905.0	2,646.0	7,905.0
<b>Artist event count</b>	<i>Mean</i>	4.0	2.8	4.6
	<i>Median</i>	2.0	2.0	2.0
	<i>Sd</i>	8.7	3.7	10.4
	<i>Min</i>	0.0	0.0	0.0
	<i>Max</i>	122.0	41.0	122.0
<b>Artist follower count</b>	<i>Mean</i>	508.4	178.7	687.2
	<i>Median</i>	71.0	61.0	78.0
	<i>Sd</i>	3,409.2	631.4	4,197.2
	<i>Min</i>	0.0	0.0	0.0
	<i>Max</i>	75,292.0	18,789.0	75,292.0
<b>Artist gallery count</b>	<i>Mean</i>	7.2	4.2	8.7
	<i>Median</i>	3.0	2.0	3.0
	<i>Sd</i>	19.3	7.1	23.2
	<i>Min</i>	1.0	1.0	1.0
	<i>Max</i>	426.0	154.0	426.0

Artwork n = 102,908 (female n = 36,177 ; male n = 66,731);  
Artist n = 11,253 (female n = 4,045 ; male n = 7,208);  
Gallery n = 1,998

**Table 4: Zero-order correlation of variables used in regression analysis**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>(1) Listing price (USD)</b>	1.00	0.07	-0.07	0.15	0.12	0.09	0.10	-0.01
<b>(2) Size (sq/cu in)</b>	0.07	1.00	-0.01	0.00	0.00	0.00	0.00	0.00
<b>(3) Year</b>	-0.07	-0.01	1.00	-0.15	-0.08	-0.09	-0.11	0.13
<b>(4) Artist gallery association</b>	0.15	0.00	-0.15	1.00	0.86	0.88	0.76	-0.02
<b>(5) Artist event count</b>	0.12	0.00	-0.08	0.86	1.00	0.85	0.69	-0.01
<b>(6) Artist artwork count</b>	0.09	0.00	-0.09	0.88	0.85	1.00	0.81	0.01

<b>(7) Artist follower count</b>	0.10	0.00	-0.11	0.76	0.69	0.81	1.00	0.00
<b>(8) Number of features</b>	-0.01	0.00	0.13	-0.02	-0.01	0.01	0.00	1.00
N artwork = 102,908								

### 3.6.3.2 Log linear mixed effect regression analysis

Results of the log linear regression with random effects for artists and gallery dummies (i.e. gallery fixed effects) are described in Table 5. Before including any predictors, a null model indicates an intra-class correlation (intercept variance to total variance) of 62.47. This suggests that a substantial amount of variance in listing prices exists between artists. Model 1 includes characteristics of the artworks and professional attributes of artists, as well as covariate balancing propensity score weights that adjust for artistic features, as described previously. This model provides a largely material and professional assessment of variation in listing prices. Results indicate that artwork size, and the artist's number of gallery associations, followers, listed works, and events are linearly related to the price of art, as expected. Given the logged nature of these variables and of the outcome, coefficients can be interpreted elastically as follows: a one unit increase in the variable is equivalent to an approximately  $\beta$  percent change in artwork listing price.

Model 2 includes artist's gender, thereby introducing the measure of gender-based social status and allowing for a test of hypothesis 2, that artwork by women is valued less than artwork by men. Given the binary nature of this variable, it can be

interpreted approximately as follows: being produced by a female artist is associated with a  $\beta \times 100$  percent change in artwork listing price. On average, women's art is listed 10% below men's, even after accounting for artist, artwork, and gallery controls, and after balancing differences in the type of work created. The 10% gap in listing prices approximates how much more women would earn if they garnered the same valuations of art given to men. This result supports hypothesis 2.

Notably, Model 2a adds an explicit control for two conceptual features of art that did not obtain sufficient balance between female and male artists via covariate balancing propensity scores weights (i.e. an absolute standardized mean difference  $\geq .09$  after balancing): *Feminist art* and art that deals with the topic of *Gender and sexual identity*. These features are included in Model 2a as binary indicators, coded 1 if the feature is present in the artwork, and 0 otherwise. The result of their inclusion is a better model fit and an increase in the estimated effect for gender, leading to an approximate 12% discount for art produced by women.<sup>3</sup>

---

<sup>3</sup> In Appendix D.1, analyses are repeated separately for each of five primary artwork types: painting, sculpture, photography, textile, and drawing.

Table 5: Mixed effect log linear model of gallery listing price for artworks

	Model 1			Model 2			Model 2a		
	Est	Std Err	p	Est	Std Err	p	Est	Std Err	p
Female				-		**	-		**
				0.100	0.015	*	0.119	0.019	*
<b>Artwork controls</b>									
Size (sq/cu in.)`	0.461	0.001	**			**			**
Year	0.001	0.001	**	0.461	0.001	*	0.463	0.001	*
Price hidden on platform			**	0.002	0.001	**	0.002	0.001	*
Gender and sexual identity content	0.115	0.006	*			**			**
Feminist art				0.115	0.006	*	0.129	0.006	*
							-		**
							0.045	0.015	*
									**
							0.129	0.006	*
<b>Artist controls</b>									
Gallery associations`	0.286	0.016	**			**			**
Events`	0.068	0.012	**	0.283	0.016	*	0.301	0.019	*
Artworks listed`	-	0.010	*	0.067	0.012	*	0.105	0.019	*
Followers on platform`	0.100	0.009	**	-		**	-		**
	0.167	0.009	*	0.101	0.010	*	0.115	0.013	*
						**			**
				0.166	0.009	*	0.158	0.011	*
Intercept	1.30	1.49		1.28	1.49		0.85	1.71	
Between artist variance		0.44			0.44			0.45	
Within artist variance		0.12			0.12			0.12	
BIC		183941.90			183917.10			146635.20	
N artists		11,253			11,253			11,253	
N artworks		102,908			102,908			102,908	
Effective sample size		53,473			53,473			53,473	

Models include covariate balancing propensity score weights and dummies for 1,998 galleries representing art in this sample.

` log-log (elasticity) interpretation: 1% change in variable  $\approx$   $\beta$ % change in listing price.

A quantile regression, which considers price segments of the market separately, provides further detail on these results (Koenker & Bassett, 1978). As seen in Table 7, the penalty for artwork by female artists increases in the top listing prices of the market. At the 30th quantile, there is an observed 7% discount for artworks by female artists, which

increases to 17% at the 90th quantile, and to 32% at the very top, 99.9<sup>th</sup> price quantile. Notably, these quantile regressions include neither artist random effects nor covariate balancing propensity score weights.

**Table 6: Quantile regression of log gallery listing price for artworks**

	Log of listing price				
	<i>q30</i>	<i>q60</i>	<i>q90</i>	<i>q99</i>	<i>q99.9</i>
<b>Female</b>	<b>-0.07</b>	<b>-0.08</b>	<b>-0.10</b>	<b>-0.17</b>	<b>-0.32</b>
Size (sq/cu in.)	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
Price hidden on platform	Yes	Yes	Yes	Yes	Yes
Art dealing with Gender and Sexual Identity	Yes	Yes	Yes	Yes	Yes
Feminist art	Yes	Yes	Yes	Yes	Yes
Gallery associations	Yes	Yes	Yes	Yes	Yes
Events	Yes	Yes	Yes	Yes	Yes
Artworks listed	Yes	Yes	Yes	Yes	Yes
Followers on platform	Yes	Yes	Yes	Yes	Yes
<i>Gallery fixed effects</i>	Yes	Yes	Yes	Yes	Yes
<i>Artist random effects</i>	No	No	No	No	No
Artwork n = 102,908					
Model does not include covariate balancing propensity score weights.					

Tables 7 through 9 contextualize this progressive discount observed by quantile regression. Table 7 conveys the proportion of market value in each bracket of the price distribution. Art by women comprises a minority share of the overall market value, at 24.25%. Though this is much higher than the auction market's 3.4% of market value held by women (Bocart et al., 2018), it is well below women's share of primary market volume, which is 35.16%. In other words, women comprise a larger portion of the market by volume than they do by value.

Table 8 dives further into the market by volume, depicting the percentages of artwork by women and men at every quantile of the price value distribution. As prices increase, women’s share of the market is reduced. This is again reflected in Table 9, which depicts the percentages of artists themselves—as opposed to the volume of their work. The last two columns of this table report the proportion of female/male artists relative to the total number of female/male artists in each bracket. A higher share of women are located in the lower tail of the value distribution, compared to the comparable share of men. At the 99<sup>th</sup> quantile of the price value distribution, where 61% of total market value is accrued, only 15.20% of female artists can be found, as opposed to 24.28% of male artists. In other words, women are more concentrated at the bottom of the market than men, and the superstar effect, in which a small number of actors comprise the majority of the market value (Rosen, 1981), is truer for the female sub-market than for the male. Combined, these descriptive statistics further depict a market in which women artists are underrepresented, less integrated into high price brackets, and undervalued across all price brackets.

**Table 7: Share of market value and volume, by quintiles of price and gender**

Quantile	Lower price threshold	% market value in bracket	female % market value in bracket	cumulative % of market value	cumulative female % of market value
> 99.9 <sup>th</sup>	\$648,106	11.38	9.16	11.38	9.16
99.9 <sup>th</sup>	\$150,000	16.44	21.19	27.82	16.27
99 <sup>th</sup>	\$25,000	32.84	21.11	60.66	18.89
90 <sup>th</sup>	\$6,800	27.10	30.50	87.77	22.48

60 <sup>th</sup>	\$2,600	9.26	35.75	97.03	23.74
30 <sup>th</sup>	\$10.40	2.97	40.74	100.00	24.25

**Table 8: Count and share of artworks, by quintiles of listing price and gender**

Quantile	Lower listing price threshold	# artworks in bracket	cumulative % of artworks	female % of artworks in bracket
> 99.9 <sup>th</sup>	\$648,106	102	0.10	9.80
99.9 <sup>th</sup>	\$150,000	855	0.93	21.52
99 <sup>th</sup>	\$25,000	9,325	9.99	22.73
90 <sup>th</sup>	\$6,800	30,776	39.90	31.26
60 <sup>th</sup>	\$2,600	30,763	69.79	36.19
30 <sup>th</sup>	\$10.40	31,086	100.00	42.17

**Table 9: Count and share of artists, by quintiles of listing price and gender**

Quantile	Lower listing price threshold	# artists in bracket	# female % artists in bracket	% of female artists in bracket	% of male artists in bracket
> 99.9 <sup>th</sup>	\$648,106	45	0.11	0.12	0.55
99.9 <sup>th</sup>	\$150,000	289	0.20	1.41	3.22
99 <sup>th</sup>	\$25,000	2,288	0.26	14.98	23.34
90 <sup>th</sup>	\$6,800	5,961	0.33	47.99	55.77
60 <sup>th</sup>	\$2,600	6,204	0.37	56.74	54.23
30 <sup>th</sup>	\$10.40	5,332	0.41	53.57	43.91

## 6.2.4 Discussion

The analyses of Part 1 and Part 2 in this study addressed the following hypotheses: that women and men make art with different characteristics, and that—net

of characteristics—art produced by women is valued less than art produced by men. Based on a combination of machine learning and multilevel regression, support is circumstantially withheld from the first hypothesis and robustly provided to the second. Specifically, random forest classification applied to 255,887 contemporary artworks on the basis of 594 aesthetic, conceptual, and material features was unable to achieve high accuracy in distinguishing between the artwork of women and men. Given comprehensive prior evidence that such methods can uncover complex patterns in equally large and complicated datasets, it is a reasonable conclusion that these analyses provide the best estimate to date on whether women and men make art with substantially different characteristics, and results indicate that they do not.

In the context of the present manuscript, which motivates analyses of gender in the contemporary art market by analogy to studies of gender in the labor market, the substantive interpretation of this is that segregation—or the creation of different types of products in the market—is unlikely to be the leading driver of gender disparity in the primary market for contemporary art. To a certain extent, this finding is anticipated. The contemporary art market is one in which a highly diverse and evolving set of products compete for attention and evaluation among a more or less unified set of consumers. While a certain subset of conceptual features related to the lived experiences of women and men are expected and were found to differ in their frequency of use between female and male artists, as a collective feature use is relatively balanced. The same is true for



categories and material features of art. Some, such as textile and graffiti arts, are conspicuously more prevalent in the artwork of women or men, but the majority are comparable in their use.

In a market as diverse as that for contemporary art, the idea that artistic categories and features alone would determine value is unlikely. Prior literature on product markets suggests that in settings of high uncertainty about product quality, actors tend to rely in non-market or product-extraneous features to organize and justify value assessments. The second set of analyses in this study thus investigated the role of gender as one such organizing feature. Conceptualizing gender as a diffuse social status characteristic that shapes judgments and downwardly biases expectations of women, I argued that it is effectively salient in the market for contemporary art, and thereby that the market cannot be understood without considering gender's influence.

Results of a log linear mixed effect regression indicated that, overall, women's art is valued between 10 to 12 percent less than men's. This is true even when relying on covariate balancing propensity score weights to adjust for imbalance in the use of artistic features between women and men. Moreover, a quantile investigation reveals an incline in which the discount for women's art grows at higher price points in the market. At highest, 99.9<sup>th</sup> percentile, women comprise only 9% of artists and there exists a 32% discount for their work. Similar findings, in which higher echelons of a market present greater gender disparity, have been found in other occupations (Gorman & Kmec, 2009)

and in the secondary auction market for art (Bocart et al., 2018). The auction market, however, is known to be reserved for artists of the past and those advanced or quickly advancing in their career. By contrast, the primary art market represents a much broader swath of artist working today, and a much higher proportion—though still significant minority—of women. Present analyses represent the first comprehensive analysis of the primary art market and the most complete accounting of product categories and features in the analysis of value differentials and the gender gap. This study is thus able to analytically distinguish between compositional effects of inequality, which are due to differences in the goods made by producers, and bias of valuation effects, which are due to differences in judgement on what is made.

An important limitation of the present study is that the data include only online listing prices. As such, they represent a subset of artists and artworks in the contemporary market. Namely, artworks listed online by gallery-represented artists. While these conditions are intentional to the scope of the present study—which aims to understand gender disparities between artists active in the primary art market today (see Appendix A)—they do limit the generalizability of findings. There is a possibility, though little reason to suspect, that art produced by non-gallery-represented artists, or art produced by galleries not represented on the Artsy platform, would present greater distinction by gender in the features of their art, or less disparity in the prices of art by women and men. It is also possible that, due to privacy and publicity of high-net-worth

individuals and their purchases, the highest primary art sales are not listed online. If so, the present results may obscure an effect in which the truly highest price bracket reveals less gender disparity. Such may be the case if, for instance, the scarcity of female artists in these upper echelons of the market results in a premium for their work. Some evidence for this is found in specific portions of the secondary auction market (Bocart et al., 2018), though those results relate to historical artworks. Furthermore, while online sales presently represent a minority of primary market art sales by galleries (25%), that share is rapidly growing, aided by sites like Artsy and necessitated in 2020 with the necessities of the COVID-19 pandemic (*The Art Market 2021*, 2021).

It should also be noted that several of the factors controlled for in the present regression analyses of part 2 in this study—including the number of galleries with which an artist associates and their popularity among online viewers—are likely to absorb a degree of the gender disparity that explains the gap in prices for women and men’s art. For example, the number of galleries with which an artist associates is a statistically significant predictor of listing price, with more gallery associations correlating with higher prices. On average, women artists are associated with only two galleries, however, while male artists are associated with three. Thus, controlling for this factor obscures the potential gender bias that leads female artists to associate with fewer galleries, which could impact listing prices. Furthermore, as women now comprise the majority of fine art degrees, and yet only a third of gallery represented artists (*The Art*

*Market 2019, 2019*), present analyses are likely to suffer from so-called survivorship bias in which those who meet requirements to be included in analysis are statistically different, in substantively meaningful ways, from those who do not. In sum, without a better understanding of the occupational patterns that lead women and men to gallery representation, I must conclude that analyses here are but a partial and likely conservative picture of gender disparity in the primary art market.

Finally, the outcome predicted in this study is gallery listing price. Sales price was not available, because a significant portion of sales did not, at the time of this study, take place through the Artsy website. While prices were commonly listed on the site, or were available in metadata, sales primarily transacted directly between buyers and the gallery representing the artwork on the Artsy site. Still, no previous scientific study has had access even to listing price information in the primary art market. Historically scattered across thousands of independent galleries, the rise of online platforms that centralize the listing of art has created new datasets that provide unprecedented opportunities to investigate inequality in a market that includes thousands of creative professionals worldwide. Further research would be required to investigate the relationship between listing and sale prices. A common norm in the primary art market, for example, is to provide discounts at point of sale. Whether such discounts mitigate or exacerbate the types of gender inequality observed here remains to be studied.

In the most complete qualitative analysis of contemporary art pricing, Velthuis (2005) explains that artworks are not priced individually, but instead follow norms that communicate information about the reputation of artists, the social status of galleries and dealers, and the quality of the artworks that are traded. Artists play their part in this process, and so present results do not provide conclusive evidence of bias on behalf of galleries and consumers. Rather, more specifically, they provide evidence of a gap in gallery listing prices that disadvantage female artists. A study of the role played by artists themselves in determining listing prices must be carried out for further understanding, but is beyond the purview of present data and analyses. In addition to the potential for female artists to undervalue their quality of work relative to men, as has been found in other fields and professions (Hewitt, 2015; Huang, 2013; Ortega et al., 2015), artists and galleries are likely to consider buyer preferences and historical market demand in their valuation of a work, both of which may result in sale-optimized prices that nevertheless devalue the labor of female artists.

To the degree that the gender disparities observed in this study are due to gallery and market pressures, and not to self-evaluations of female and male artists, practical means to closing the gender gap are likely to come from a change in the transparency of listing prices across the market and within galleries. Such transition has already begun, with the rise of online platforms that list art for sale. Greater transparency in listing prices would allow for greater scrutiny and accountability of disparities not only by

gender, but also by other potentially salient characteristic of producers, such as their race or ethnicity. In other creative professions, blinding audiences and judges to the identity of the artist has been shown to reduce gender disparity in assessment and to significantly improve gender representation in career advancement (Goldin & Rouse, 2000). Blinding audiences to the gender of artist in the contemporary art market is less feasible, however, as the uniqueness of products is often inextricable from the identity of individual creator. The individual creator also often maintains a personal relationship with galleries and collectors. As such, transparency in pricing, as well as institutional policies, is the more reasonable way to reduce disparity in the economic outcomes of female and male artists.

## 4. Study B

The notion of a gender gap in the labor market is commonly used to describe wage, hiring, and status disparities between working women and men. While the concept has circulated broadly into academic and popular parlance, however, the notion of the gender gap refers to several realities. In the previous section, I addressed two common attributions of the gender wage gap in the labor market at large: that women select into different occupations, and that women's labor is valued less even when fulfilling the same occupation. A third line of reasoning combines the prior two. Known as devaluation theory, this perspective proposes that, beyond women selecting different occupations or being paid differently for the same occupation, gender also influences wage structure because jobs and tasks filled mostly by women pay less than they would if they were filled mostly by men (England, 1992, 2005; Heilman, 2001; Steinberg, 2001). According to devaluation theory, the primary mechanism for this process of devaluation is one in which economic compensation is biased by sexist, institutionalized beliefs about the worth of culture and tasks associated with women. Plainly, the social status of women propagates onto the tasks and titles of labor women typically perform, such that those tasks and titles are valued less than the male equivalent.

Tests of devaluation theory have attended to the ways in which gender composition across occupations impacts material compensation (i.e. wages), symbolic compensation (i.e. prestige), and the relationship between the two (Magnusson, 2010).

The resulting literature is a landscape of robust but contradictory findings that leave substantial room for further research. Moreover, while a theoretical lineage within cultural sociology preempts tests of devaluation theory in the context of product markets—where the social status of women may devalue not only the occupations they typically fill, but also the products or features of products they typically produce—most empirical analysis has been carried out in labor markets.

The following study contributes to the literature on gender, markets, and inequality by expanding the evidence base of devaluation theory, and by presenting the usefulness of this cultural theory of labor to the study of cultural product markets. More specifically, I build on my previous analyses of the gender gap in gallery listing prices for contemporary art to account for the effective salience of gender in the valuation of artistic features and categories.

#### ***4.1 Material and economic evidence of devaluation theory***

The most common outcomes in prior studies of devaluation theory are material. They address, for instance, wage differentials between occupations filled predominantly by women and those filled predominantly by men. Even after accounting for additional factors that may impact compensation, such as industry and required levels of education, training, experience, physical hardship, or risk, research in this area has persistently found that both women and men earn less in “female” occupations (Alksnis et al., 2008; Cohen & Huffman, 2007; England, 1992; England et al., 1988, 1994; Johnson &



Solon, 1986; Macpherson & Hirsch, 1995; Tam, 1997), with the percent of female workers in an occupation being negatively associated with wages (Baron & Newman, 1990; Dinovitzer et al., 2009; England et al., 1994, 2007; Huffman & Velasco, 1997; Kilbourne et al., 1994; Levanon et al., 2009; Peterson, 1992). This is particularly true in markets that have ambiguous valuation criteria (Baron & Newman, 1990). Furthermore, even within occupations, job tasks typically associated with women tend to be valued less than those typically associated with men (England, 1992; Guy & Newman, 2004; Hall, 1993; Hughes, 1945). Abbott (2010) describes this pattern as a sort of fractalization, in which gender segregation regenerates itself in new contexts; for example, in highly selective occupations like medicine and law, the influx of female workers has coincided with female enclaves and specialties that tend to be compensated less (Cohen et al., 2009; Ku, 2011; Levanon et al., 2009). Rather conclusively then, studies considering the material compensation of comparable jobs held mostly by women or men tend to find support for devaluation theory.

In the effort to test devaluation theory, the burden of scholars rests in the task of demonstrating that jobs filled predominantly by women and men are otherwise comparable. If, for example, jobs filled by women require less training, less risk, or less risky work environments, then the fact that they pay less may be rationalized. Most studies address this issue of comparability by controlling for factors other than gender composition that may reasonably justify compensation. This leaves the proportion of

women in an occupation to demonstrate an independent effect on wages. Other studies have relied on longitudinal data that allows for investigating how changes in the gender composition of occupations over time impacts their economic worth (England et al., 2007). In both approaches, support for the material devaluation of predominantly female occupations has been found.

In the present research, I analyze a product market, rather than a labor market. Specifically, I analyze the product market for art. Doing so assuages several of the concerns generally associated with establishing comparability between the type of work performed by women and men. While certain forms of art do indeed require more training, time, or physical resources to create, we also know that women now comprise the majority of graduates with fine art degrees and, as demonstrated in the previous study of this dissertation, often segregate from men on purely conceptual or thematic features of their art. As such, a test of devaluation theory in this market is arguably more tenable, because the laborers and their work are more comparable than across occupations in the labor market more generally.

In devaluation theory, status beliefs about the relative respect, competence, and worth of women downwardly bias judgments of the work and of the type of work women typically perform (Correll & Ridgeway, 2006; Tomaskovic-Devey, 1993). It is a reasonable proposition that those same status beliefs may influence judgments of the

products and qualities of products that women typically create. Assessing if this is so is the aim of the present study.

#### ***4.2 Symbolic evidence of devaluation theory***

As a rebuttal of the evidence provided in support of the material devaluation of women's work in the labor force, Oschenfeld (2014) argues that men are more likely to select into "breadwinner" occupations than women. That is, net of education and training, occupations selected by men tend to be culturally understood as career-driven and are provided with high economic return. While the author concludes that this finding refutes devaluation theory, the question that remains is obvious: why, net of practical justifications for material worth such as experience and education, are predominantly male occupations perceived as breadwinner roles and rewarded with higher compensation? Such a trajectory of thought turns to the symbolic form and implications of devaluation theory.

The symbolic worth of occupations, often measured as occupational prestige, reflects the moral standing and reception of honor they hold in a society. Symbolic worth is an alternative to material compensation when considering the relative value of work. In the literature on devaluation theory, however, studies of symbolic worth are both less settled and less common than their material counterparts. Some scholars have found that occupations filled predominantly by women do not differ in their symbolic worth, or prestige, when compared to occupations filled predominantly by men (Crino

et al., 1983; England, 1979; Suchner, 1979; M. C. White et al., 1981). Others have found that predominantly female occupations are either more (Baunach, 2002; Rosenfeld, 1980) or less (Bose & Rossi, 1983; Touhey, 1974a, 1974b; Xu & Leffler, 1992) prestigious than predominantly male occupations. And still others have found that it is the least gender segregated occupations that garner the highest (Magnusson, 2009) or lowest (Valentino, 2020) judgments of symbolic worth. In other words, the literature is inconclusive.

The present study does not attend to judgements of symbolic worth. It remains focused on material, or economic, valuation. In the analyses that follow, I investigate how women's and men's proportional use of art features impact disparity in listing prices for contemporary art. Nevertheless, I highlight studies of symbolic worth because the market I consider is deeply symbolic. In art, uniquely more than other in commodities, economic value is inextricable from symbolic meaning. Art is, "a two-faced reality, a commodity and a symbolic object" (Bourdieu, 1993). As such, the inconclusive literature on the symbolic worth of gender in occupational characteristics elicits hesitancy: will aesthetic, conceptual, and material features associated with the work of women artists diminish or exacerbate the economic devaluation of women's labor, as was observed in the previous study? Indeed, even as I have argued that the symbolic nature of art eases certain empirical requirements when establishing the comparability of women's and men's work, it also obfuscates a clear expectation of whether and how gender will relate to economic value. For reasons discussed further

below, I hypothesize that characteristics of art predominantly associated with the work of women will be devalued relative to those predominantly associated with the work of men, but it remains true that the inconclusiveness of prior research on the symbolic worth of gendered occupations leaves this hypothesis far from certain of gaining support.

### ***4.3 Devaluation theory in the context of product markets***

Empirical tests of devaluation theory in the context of product markets, rather than the labor market, are lacking. Application of the theory to product markets, however, is not new. For most economic sociologists, products in a market are understood as objects imbued with meaning by the social world around them (Fligstein & Dauter, 2007). It is therefore not a stretch to imagine that the meanings ascribed by society to products may reflect the social world, and the identity of producers and consumers within it. This can be true of goods in any market, but is perhaps nowhere truer than in markets for art, which have a special tie to identity.<sup>4</sup>

In the sociological literature on cultural consumption, a primary tenant is that there exist structured connections between the social categories of consumers and the categories of products they choose to consume, as well as the meaning and value of those products (Bourdieu, 1979, 1983; DiMaggio, 1987; Lizardo & Skiles, 2016). At least

---

<sup>4</sup>. This is similar to so-called “peculiar” goods, which are goods not easily valued by monetary prices without violating social norms, such as nature or human life (Fourcade, 2011).

as far back as Thorstein Veblen's notion of conspicuous consumption (1899), this tenant has flourished and is now repeated in the common finding that consumer preferences tend to align personal identities with the perceived identity of product audiences—be those by gender, race, class, or region (Bryson, 1996, 1997; Halle, 1996; Lizardo & Skiles, 2016). Heavy metal, as shown by Bryson (1996), is a genre of music seen as “male” due to the social demographic characteristics of those who consume it. Likewise abstract art, indigenous sculpture, mass-produced sunsets, and vivid religious representations have patterned social demographic identities, at least within the bounds of New York City, as shown by Halle (1996). The cultural omnivore literature—both its proponents and its dissenters—is an exemplary strain of this intellectual agenda attributing a duality between categories of cultural goods and the social categories of their consumers (Goldberg, 2011; Peterson, 1992; Peterson & Anand, 2004; Peterson & Kern, 1996).

Consumers are not the only identities to influence product meanings however. Indeed, the broader theories that underly studies of consumption are also acutely aware of the relationship between products and producers—especially in the case of cultural products. Bourdieu, for example, theorized not only that the audiences who choose to consume a product shape its meaning and value, but also that the characteristics of producers—and particularly the forms of social capital ascribed to and gained by them outside the field of production—bear sway (Bourdieu, 1983, 1984). That is, the qualities

of producers can propagate onto the positions they assume in the market, and from the positions onto features and products themselves (Rule & Bearman, 2015).

While generally more macro in perspective, so-called “production of culture” studies writ large attest to the fact that the cultural meaning of products derive from the social systems from which they emerge and in which they circulate (for review, see Peterson & Anand, 2004). Thus, just as the qualities of products can come to define and be defined by the audiences who consume them, so too can they define and be defined the producers who most commonly produce them. In what follows, I test this. Specifically, I test whether the social status characteristic of gender—ascribed to individual producers—adheres to qualities of art. Do aesthetic features predominantly associated with the work of women decrease the value of an artwork? And further, does there exist a process of gender congruity, whereby artists are rewarded for producing aesthetic features typical of their gender?

Social psychological work has repeatedly found that individuals are judged by gender stereotypical beliefs for behaviors in the workplace (Heilman, 2001; Heilman et al., 2004; Ko et al., 2015), as well as in education (Steele et al., 2002) and assessments of creativity (Proudfoot et al., 2015). Occupational theorists have extended this perspective to demonstrate that entire occupations can be gender-typed (Brynin & Perales, 2016; Riach & Rich, 2002; Rudman & Phelan, 2008), with women—in a version of benevolent sexism (Fiske & Glick, 1996)—being rewarded for performing labor typically perceived

as female (Guy & Newman, 2004). Whether and to what extent a similar logic applies to the production of goods is an intuitive next inquiring in the aim of understanding how gender impacts markets, but it is a form of inquiry that has only very recently been taken up in empirical research.

While not motivating analysis on the explicit basis of devaluation theory, both Tak, Correll, and Soule (2019) and Kricheli-Katz, Regev, and Correll (2019) theorize a “status belief transfer” in product markets, by which the diffuse status characteristic of gender is transferred from the producers onto the products they create, or from merchants onto the products they sell. Through a series of online experiments with simulated data, Tak et al find both that entire product markets can be gender-typed, and that there is an asymmetric negative bias, wherein products made by women in a male-typed product market (the market for craft beer) are disadvantaged, but the same is not true for products made by men in a female-typed market (the market for cupcakes). Similarly, Kricheli-Katz et al find that women are penalized more for selling products on Ebay that are perceived to be owned typically by men than are men for selling products perceived to be own typically by women. These effects are found to be amplified under conditions of price uncertainty.

I build upon these studies both theoretically and empirically by proposing not only that men’s and women’s products face differential valuations in male- and female-typic markets, but also that the features of products themselves, even within a single



market, can become harbingers of the status ascribed to producers on the basis of gender. That is, while Tak et al test for disparity between comparable products in two gender-typic markets, I test for disparity between comparable but gender-typic features in a single product market. And while Kricheli-Katz et al focus on the gender typicality of products according to those who consume them, I focus on the gender typicality of products according to those who produce them. In so doing, I rely on real world market data, thereby strengthen the evidentiary base surrounding devaluation theory, and enlightening the several means by which gender can come to bear products in a market.

#### **4.4 Hypotheses**

Drawing on the previously discussed devaluation literature of the labor market, I propose and test an analogous hypothesis in which female-typic product characteristics are valued less than male-typic product characteristics. In the context of the art market, the following hypothesis should be supported:

*Hypothesis 3: Artwork with characteristics made predominantly by female artists are devalued relative to artworks with characteristics made predominantly by male artists.*

Further, provided with prior and persistent evidence that the individuals are rewarded for their performance of gender-congruous tasks, I hypothesize that the value of art will be higher when the characteristics it depicts are more typical of the gender of the artist.

If this is true, the following hypothesis should be supported.

*Hypothesis 4: Artwork will be valued more when the artist depicts characteristics predominantly produced by artists of their gender.*

## 4.5 Data and method

To test my hypotheses, I extend the regression analyses of the previous study (see section 3.6), with the addition of a constructed measure for how “gender-typic” an artwork is, according to the categories and aesthetic, conceptual, or material features it depicts. To address the potential that female and male artists are differentially impacted by their production of gender-congruous features (hypothesis 4), I also interact this measure with the observed gender of the artist.

My operationalization of the gender-typicality of an artwork is measured via the winsorized mean of the relative gender frequency ratios for all features appearing in an artwork. Formally, this measure is calculated as follows. First, for each feature, the relative frequency ratio for female and male artists is calculated. That is, the absolute frequency of a feature’s use among female artists (the number of unique female artists using the feature) is divided by the overall size of the female corpus (the number of unique female artist-feature combinations in the entire dataset). Then, the ratio is taken between this value for female artists and the same value for male artists:

$$\gamma_i = \frac{\text{freq}(\text{female}_i)}{\sum \text{freq}(\text{female}_{i\dots k})} \bigg/ \frac{\text{freq}(\text{male}_i)}{\sum \text{freq}(\text{male}_{i\dots k})}$$

Where  $\gamma_i$  is the relative frequency ratio for feature  $i$ . Notably, this relative frequency ratio is the same measure used in Figure 3 of this manuscript to investigate

the gender differences of feature use across artists. The final measure of gender-typicality for an artwork is calculated by the mean of the winsorized relative frequency ratios for all features that the artwork depicts, according to the coding of art historian as described in the Artsy Art Genome Project previously discussed<sup>5</sup>:

$$\frac{\gamma_1 + \gamma_2 + \gamma_3 + \gamma_{n \text{ feat}}}{n \text{ feat}}$$

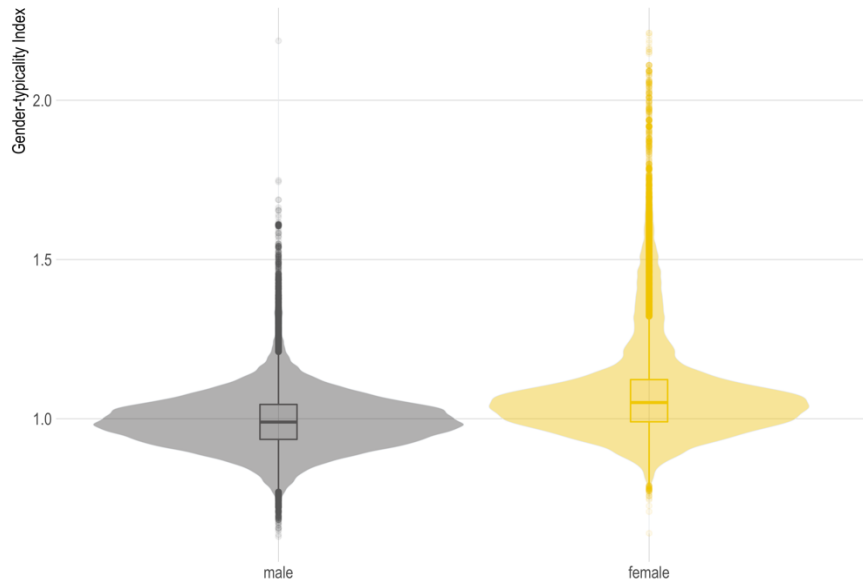
Winsorized mean is implemented because descriptive analysis of relative frequency ratios indicated the presence of extremely high outliers. While the median and third quartile of the relative frequency ratios are 0.96 and 1.16, the top three values are 5.43, 6.75, and 52.50, representing the Femininity, Knitted and Crocheted, and Feminist features respectively. Grubbs's tests further indicate that outliers are present in the higher but not the lower ranges of the distribution. As a result, the mean relative frequency ratio for this feature was winsorized to within the upper 99<sup>th</sup> percentile prior to the calculation of the final "gender-typic" variable.

Descriptive statistics for the gender-typicality measure are provided in Figure 5. Given the results of section 3.5 of this manuscript, which indicated that women and men do not make substantially different types of art, one would expect this measure to

---

<sup>5</sup> As this measure aims to compare the distribution of feature use between male and female artists, it is calculated at the level of the artist; that is, the use of each feature is counted only once for an artist, rather than once for each artwork. This diminishes the influence of any single artist's use of a feature across a large number of artworks and thereby more accurately represents the relative frequency of a feature across female and male artists.

present only modest differences between women and men. And, indeed, that is the case. A t-test indicates a statistically significant difference in means on this gender-typicality scale, with women statistically above men (i.e. producing more female-typic features, on average), as expected. However, this difference is small. As seen in Figure 5, the distributions for women and men are not highly skewed, with both clustering around 1. By in large, however, art by women present more pronounced outliers than do artworks by men on this scale.



	Min	1 <sup>st</sup> Q	Median	Mean	3 <sup>rd</sup> Q	Max	Std Dev
<i>Female</i>	0.640	0.990	1.051	1.091	1.123	2.211	0.172
<i>Male</i>	0.630	0.935	0.990	0.995	1.045	2.187	0.096
Total	0.630	0.952	1.010	1.029	1.072	2.211	0.136

**Figure 5: Artworks plotted by the gender-typicality index, and by the gender of the artist (n = 102,908)**

## 4.6 Results

To test hypothesis 3, that female-typic features of art are valued less than male-typic features of art, I examine the model fit and effect sizes of Model 3 presented in Table 10. This model maintains the log of the gallery listing price of artworks as the outcome, with the primary predictor being the standardized gender-typicality of the work. In addition, I include all predictor variables used in final analysis of the previous study: artist gender, artwork size, artwork year, number of gallery associations, number

of artwork listed on the Artsy platform, and number of followers on the platform. I also control for the number of artistic features applied to the artwork via the Art Genome Project, as well as a dummy indicator for whether the artist and gallery chose to make the artwork price public.

As the gender typicality variable has been standardized, the results of Model 3 indicate that a one standard deviation increase in gender-typicality—equivalent to .14 in the original scale—is associated with a 3% discount in listing price ( $\beta = -0.030$ ;  $p < 0.001$ ). Given that this variable is the standardized mean of the log frequency ratio for of female to male artists, an increase is associated with artwork depicting more female-typic characteristics (that is, features more often used by female artists).<sup>6</sup> These results thus provide support for hypothesis 3: artworks with more female-typic qualities are priced lower than artworks with more male-typic qualities, net of the observed gender of the artist. The observed gender of the artist remains a persistent and significant predictor of listing price, with women’s art listed at a discount of 8% relative to men’s ( $\beta = -0.084$ ;  $p < 0.001$ ). In combination, the observed gender and gender-typic coefficients provide evidence that variation in artwork listing prices is due partially to the devaluation of female labor and partially to the devaluation of characteristic that are typical of female labor.

---

<sup>6</sup> This analysis was also conducted with a polynomial version of the gender-typicality index. Model fit was very slightly worse and present theory does not require this operationalization.

**Table 10: Mixed effect log linear model predicting gallery listing price for artworks**

	Model 3			Model 4			Model 5		
	Est	Std Err	p	Est	Std Err	p	Est	Std Err	p
Female	-0.084	0.015	***	-0.085	0.015	***	-0.096	0.018	***
Gender-typicality	-0.030	0.003	***	-0.045	0.005	***	-0.060	0.006	***
Female * Gender-typicality				0.022	0.006	***	0.021	0.007	**
<b>Artwork controls</b>									
Size (sq/cu in.) <sup>`</sup>	0.459	0.001	***	0.459	0.001	***	0.459	0.001	***
Year	0.001	0.001		0.001	0.001		0.003	0.001	***
Price hidden on platform	0.115	0.006	***	0.115	0.006	***	0.129	0.006	***
Features applied <sup>`</sup>	0.120	0.009	***	0.119	0.009	***	0.102	0.010	***
Gender & Sexual Identity							0.070	0.019	***
Feminist							0.083	0.027	***
<b>Artist controls</b>									
Gallery associations <sup>`</sup>	0.285	0.016	***	0.285	0.016	***	0.295	0.019	***
Events <sup>`</sup>	0.066	0.012	***	0.066	0.012	***	0.076	0.014	***
Artworks listed <sup>`</sup>	-0.102	0.010	***	-0.102	0.010	***	-0.151	0.013	***
Followers on platform <sup>`</sup>	0.165	0.009	***	0.165	0.009	***	0.172	0.011	***
Intercept	1.96	1.49		1.91	1.49		1.43	1.68	
Between artist variance		0.44			0.44			0.44	
Within artist variance		0.12			0.12			0.13	
BIC		183669.90			183676.60			151389.40	
N artists		11,253			11,253			11,253	
N artworks		102,908			102,908			102,908	
<i>Effective sample size*</i>		53,473			53,473			53,473	

Models include both covariate balancing propensity score weights and dummy variables for 1,998 galleries.


<sup>`</sup> logged (elasticity) interpretation: 1% change in variable  $\approx \beta\%$  change in listing price. All other coefficients interpretation: 1% change in variable  $\approx 100*\beta\%$  change in listing price.

To test the final hypotheses of this chapter, that artists are rewarded for producing artistic characteristic that are typical of artists of their gender, I interact the observed gender of the artist creating each artwork with the measure of the artwork's gender-typicality. This allows the latter to have unique coefficients for female and male artists. Results demonstrate that both the coefficient estimate for gender-typicality and

the interaction term with observed gender are statistically significant, but opposite in valence. For male artists, a one standard deviation increase in the gender-typicality of an artwork—an increase in the production of more female-typic qualities—is associated with an approximate 4.5% decrease in listing price ( $-\beta = .045$ ;  $p < .001$ ). By contrast, this discount is smaller among female artists ( $\beta = .022$ ;  $p < .001$ ). For female artists, a one standard deviation increase in the gender-typicality of a work is associated with a 2.5% decrease in listing price.

The predicted values from Model 4 are plotted by gender in Figure 6, visually representing the two primary results: an overall negative slope indicates that artists of both genders experience a decrease in value for the production of more female-typic art, but the separate slopes implies that this effect differs by the gender of the artist. A smaller penalization for women producing more female-typic qualities of art results in a lack of gender gap when artwork is more female-typic. By contrast, there is a statistically significant gender gap in listing prices for artwork that is more male-typic. In other word, male-typic art is valued higher and provides men with a statistically significant boost, while female-typic art is valued lower and provides women with no equivalent boost. The histogram above the line plot indicates that the majority of artworks fall into the distribution of effects in which there exists a significant gender gap.





**Figure 6: Predicted listing price for artworks, derived from Model 4**

As a reminder, the models present here include the covariate balancing propensity score weights calculated in Study A of this manuscript. Model 5 is a robustness check that removes artworks for which gender was imputed (as opposed to being provided by the artist or gallery) and adds an explicit control for two conceptual features that did not obtain sufficient balance, defined as an absolute standardized mean

difference of .09 or greater after weighting. These features are *Feminist art* and art that deals with the topic of *Gender and sexual identity*. I add these features to the model as binary indicators, coded 1 if the feature is present in the artwork, and 0 otherwise. The result of this robustness check is very similar to the primary analysis in Model 4. However, the main effect for the gender typicality index is slightly increased, suggesting that for male artists, a one standard deviation increase in the gender-typicality of a work is associated with an approximate 6% decrease in listing price ( $\beta = -.060$ ;  $p < .001$ ). For women, the discount remains smaller by 2%, indicating an approximate 4% discount in listing price for every standard deviation on the gender-typicality index ( $\beta = .021$ ;  $p < .01$ ). There also remains a persistent, overall effect of artists gender, suggesting a discount of almost 9% ( $\beta = -.087$ ;  $p < .001$ ) for art by women at the mean of the unstandardized gender typicality index, equivalent to a value of 1.<sup>7</sup>

## **4.7 Conclusion**

This study addressed, adapted, and tested a cultural theory of labor market, devaluation theory, in the context of a cultural product market, the market for contemporary art. I hypothesized that, beyond women and men producing art with different characteristics or having their art valued differently for the same characteristics, gender also shapes the art market because features and categories of art

---

<sup>7</sup> See Appendix D.2 for further robustness check details.

produced mostly by women are valued less than they would be if they were produced mostly by men. The proposed mechanism behind this effect is one in which sexist and institutionalized beliefs about the worth and competence of women can spread from female producers onto the qualities of work they typically produce. If such a process exists, one would expect art with more female-typical qualities to be valued less than art with more male-typical qualities, independent of the gender of the artist producing the work.

Leveraging data from an online art market platform, which provides detailed information coded by art historians about the aesthetic, conceptual, and material features of art, I developed a product-level measure of gender-typicality that allows for a test of devaluation theory in the contemporary art market. Results of multilevel regression predicting the gallery listing price for 102,908 contemporary artworks indicate that, net of the observed gender of the artist, artworks are valued less if they depict features predominantly made by women. This supports the broadest hypothesis of devaluation theory, and replicates findings from the labor market in which jobs filled mostly by women are valued less than they would be if they were filled mostly by men (F. D. Blau & Kahn, 2016).

In addition to this primary effect of how the gender-typicality in artistic features impacts the economic value of art, the combined effect of gender-typicality with the observed gender of the artist provides an assessment of whether women and men are

rewarded for producing art with qualities stereotypical of their gender. Results of an interaction model indicate a gender gap in the listing prices for male-typic art, but not in the listing prices for female-typic art. Such a pattern aligns with the only two prior studies on the influence of gender in product markets, both of which found that women are penalized more for producing or selling goods that are typically associated with men (Kricheli-Katz et al., 2019; Tak et al., 2019).

#### ***4.8 Discussion***

The contributions of this study can be seen with reference to the present state of literature on how social status characteristics, and particularly gender, influence the outcomes of producers in product markets. Very little research prior to the present has explicitly focused on this question. Thus, but gleaning insights from the literature on labor markets, where consideration of this topic is more abundant, in this study I examined how the gender typicality of artistic features in the market for contemporary art relate to the valuation of artworks, net of the artist's gender. My modelling strategy thus addressed the common intuition that women and men may make different types of goods in this market, while also attending to the potential rebuttal: so what if they did?

The analysis presented in this study is one of only three recent efforts to extend our understanding of gender as a social status characteristic to the context of product markets, and the first to do so using real world market data on producers and the products they create. Broadly, my findings imply that the majority of artworks by both

women and men do not portray highly gender-differentiated qualities, but that there is nevertheless a relationship and an independence between the effect of an artist's gender and the effect of producing gender-typical features on the price of artworks today. Specifically, artworks are valued less if they are produced by a woman, and less again if they depict features of art typically produced by women.

These findings clearly bolster a case for the integration of social status characteristics into market research. They challenge the idea of what constitutes a market structure and encourage scholars to attend to the social systems and frames of cultural belief in which markets are embedded and entangled. Past studies of social status in product markets have chiefly focused on organizations, rather than individuals, and operationalized status on the basis of associations and events taking place in the market. While prolific and useful, this work has sidelined the study and theoretical development of individual producers and status ascribed to them outside of the market, on the basis of sociodemographic and other identity characteristics. In cultural sociology, research on the sociodemographic characteristics of audiences has flourished—for instance, in the literature on cultural omnivorousness—but again, attention to producers has been brief. The present work thus provides an important link in the literature. It raises the role of individual producers in product markets (particularly artists to the sociological study of art markets) and calls for attention to the ways in which producer identity might shape market structures of production, meaning,

and value. It also readies the field for future work on the interplay of the status characteristics held by producers, consumers, and their products.

The lack of accounting for consumer preferences is a shortcoming of this study. To the degree that the pricing norms of galleries reflect consumer preferences, engagement with the social and cognitive processes of consumers will be essential to complete the picture on why a gender gap in listing prices persists (Beckert, 2019). Bryson's (1996) symbolic exclusion model, for example, argues that preferences for symbolic goods are derived from the social boundaries of those who consume them—be those based on gender, race, class, or region. Under this rubric, genres, features, and categories associated with “negative” social group will be rejected, where “negative” is defined relative to the social attributes and understandings of the consumer (Bourdieu 1984; DiMaggio 1997; Goldberg 2011; Martin 2003), as well as the perceived attributes of the products (Lizardo & Skiles, 2016).

How might this influence the gender gap of the contemporary art market? While recent surveys indicate that women comprise a growing number of the high-net-worth individuals who may purchase art, and a majority of citizens who visit art institutions, at least the United States and Britain (Christin, 2012; DiMaggio & Mukhtar, 2004; Katz-Gerro & Sullivan, 2002), they also indicate that women remain a minority (33%) of art collectors (*The Art Market 2019*, 2019). As a result, the male perspective on art may have a unique influence on judgments of consumer preferences, and thereby on the pricing

norms of galleries and artists. Future research should investigate consumer demographics and preferences, as well as the relationship of those patterns to pricing behavior in the market for contemporary art. In many markets, the arrival of online platforms has generated data to investigate these dynamics. Employing such data in scientific research would bring greater fullness to our understanding of (in)stability and (in)equality in the art market.

Finally, the findings of this study encourage scholars of culture and markets to detangle the devaluation of labor performed by women from the devaluation of qualities of labor associated with women. While causality cannot be established by the methods of this manuscript, the theory proposed and tested argues that the social status of women is transcribed onto the qualities of art they typically create, and that these qualities then become harbingers of bias that can perpetuate inequalities. In a certain sense, this can be thought of as translating gender bias into a femininity bias, creating an opportunity for a future in which women and men's labor is valued equally, but traits of femininity are still valued less than traits of masculinity. Such an insight underscores the need for cultural sociology in the study of markets, as there are substantial theoretical foundations on which market studies could build a better conceptualization how status hierarchies, cultural symbols, and social system might interact to shape economic outcomes (Bourdieu, 1983; Peterson & Anand, 2004; Simmel, 1957; Veblen, 1899).

As a rare test of devaluation theory in the context of a product market, this study primes future research into how and under what conditions cultural beliefs about gender spread beyond the bodies of women and men and into cultural symbols to (re)produce gender disparity in economic spheres. When studying economic outcomes, market studies must engage with the systems and hierarchies of the societies in which markets are embedded. This research insists that it is not enough to assume that producers in a market are judged solely on the qualities of their goods and not also by the qualities of their person, nor is it enough to assume that the notion of quality itself is an unbiased assessment of worth. The processes of human, symbolic, and economic valuation are entwined, and without considering the lot, markets are likely to remain misunderstood in research, and inequitable in reality.



## 5. Conclusion

The opening of this dissertation presented statistics that document long-standing gender inequality in the art world. Both the theory and the empirical analyses that followed were aimed at explaining the persistence of this inequality today. Using an unprecedented dataset derived from an online marketplace, and combining theories of economic, gender, and cultural sociology, the three studies laid out in this dissertation address supply, demand, and devaluation theories for the gender gap in listing prices of contemporary art market. Results challenge traditional notions of what constitutes a market structure and improve our understanding of how social status characteristics come to bear on product market outcomes. The present dissertation thus makes a clear contribution to the sociological study of markets by focusing on how social systems interplay with what is traded ([Aspers, 2008](#)). It also makes contributions to the study of culture and gender by demonstrating how categories and features of cultural products can inherit and become harbingers of the systemic biases held toward social groups.

The first study of this manuscript addressed a prominent and colloquial supply-side argument of the gender gap in the art world, which is that women and men produce art with different characteristics. I tested this hypothesis using a combination of machine learning classification methods and descriptive statistics. Results indicate that while women and men do differ in their use of certain aesthetic, conceptual, and material features of art, the infrequency of these features and the lack of substantial

differentiation in most features implies that segregation in the type of artwork created by women and men is unlikely to be the leading cause of disparity in their listing prices.

The second study of this manuscript addressed a demand-side argument of the gender gap in listing prices. Using multilevel logistic regression and covariate balancing propensity score weights to account for artistic features, I tested the hypothesis that gender—as a diffuse social status characteristic of producers—can influence judgments of women’s and men’s labor. Results provide support for this hypothesis. Artworks by women are valued less than artworks by men, even when they depict similar aesthetic, conceptual, and material features. Specifically, women’s art is discounted at approximately 10% of men’s, which rises to over 30% at top price bracket of the market. These findings align with studies from other industries and across the labor market, where women also face greater compensation disparity as they advance in careers (Albrecht et al., 2003; Magnusson, 2010).

Finally, the third study of this manuscript addressed the potential that not only the labor of women is devalued in market for contemporary art, but also the qualities predominantly associated with their art. Specifically, I hypothesized that aesthetic, conceptual, and material features predominately made by female artists would be valued less than those predominately made by male artists, net of the gender of the artist producing them. Findings, again from multilevel logistic regression analysis, support this hypothesis. Further, they support the hypothesis that women and men are

rewarded for producing art with qualities that are more typically produced by artists of their same gender.

What do these findings mean for the study of gender—and social status characteristics more generally—in product market research? There are two primary implications: first in relation to how economic sociologists operationalize the concept of status, and second in relation to how both economic and cultural sociologists understand the relationship between the concepts of product quality and producer status. With regard to how economic sociologists operationalize status, the considerable variation in art price by artist gender found in this dissertation is indicative of the need for future research to address not only the status of producing organizations—generally measured by associations and events taking place in the market (Podolny, 1993a)—but also the status of producing individuals, as measured by the sociodemographic and other identity characteristics that those producers bring to market. In particular, the interaction of these two forms of status is likely to yield many insights into market dynamics.

Beyond gender, other social characteristic—like race, ethnicity, and class—are likely to be salient enough in many product markets to shape processes of valuation. These traits thus present market scholars with the complicated endeavor of understanding intersectional space in theory and market function. I propose that the breadth and importance of this task is deserving of a new emphasis within economic

sociology. Only by understanding the how producer identities shape market outcomes can the dynamics of markets, and particularly their inequalities, be fully understood.

The second implication of this study to academic research relates to the how economic and cultural sociologist understand the relationship between product quality and producer status in markets. Quality, as described earlier in this manuscript, is generally understood to be a feature of products or services. Precisely, quality is the attributes of products and labor that are considered desirable in a market (Gould 2002, p1153; Olliver 2004, p. 148). Status, on the other hand, is a social position of individuals and organizations based on honor and rewarded with privileges (Weber, 1978, p. 932). It is a common assumption in market research that producer status acts as a signal for the underlying quality of products, especially when that quality is uncertain or hard to judge (Lynn et al., 2009; Spence, 1974). The findings of this dissertation present a theoretical refinement for this assumption.

When status is operationalized on the basis of social ties within the market, its use as a signal for quality makes sense. In theory, it remains distinct from the meaning of the qualities themselves. However, when status is operationalized on the basis of sociodemographic and other identity characteristics, this relationship between status and quality is complicated. As the findings of the third chapter in this dissertation attest, the underlying qualities of products can inherit meaningful bias targeted at their producers. In this light, what is there to distinguish between the desired attributes of

products (their quality) and the status or social position of producers? This is a question of economic theory ready for sociological engagement. Current theory assumes that quality and status are decoupled in contexts of uncertainty (Lynn et al., 2009), but perhaps instead they are fused. Scholarship must reckon with this conundrum, as it's empirical implications for the study of markets and market inequality could be broad.

This dissertation has special relevance to the study of art and art markets. It is the first study to investigate prices the primary market for contemporary art, and the most detailed assessment of gender inequality in this space. As such, it provides empirical evidence for the longstanding question of why there have been so few great female artists (Nochlin, 1971), and why female artists remain at a disadvantage today. In this, it is both performs and challenges Bourdieu's injunction to ask, "not what the artist creates, but who creates the artist" (Bourdieu, 1993, p. 147). This research affirms that such a distinction is difficult, if not impossible, to achieve. However, it also affirms that the aesthetic qualities of art alone are not enough to understand artistic greatness. Inequality research in the sociology of art is rare, but efforts to understand genres, features, meaning, and value will be inhibited without it. Scholars of art must engage with the status hierarchies and social systems from which artists emerge, in which they are received, and by which they communicate their work. Prior qualitative research on the valuation of art has documented social norms of pricing that deviate from standard

economic logic (Velthuis, 2003). The present research contributes to this by demonstrating that gender too is salient structure of art market valuation.

What practical implications do the findings of this study have for the art world?

The evidence that the gender gap in gallery listing prices is driven both by the devaluation of female labor and by a devaluation of qualities of art predominantly made by female artists suggests a conundrum. Of course the task remains of valuing women's and men's labor commensurably when they produce the same qualities of art; but, the findings here suggest that a gap will remain so long as there exists gender differences in the rate at which various qualities are produced. If the qualities of art that differentiate women and men are also different in the effort or material cost expended to create them, then some gender gap may be rationalized. However, to the degree that the aesthetic, conceptual, and material qualities of art produced by women and men are comparable, even if they are different, a gender gap should not remain. If it does, the market is faced with two solutions for narrowing the remaining gap.

First, theoretically, female and male artists could be encouraged to produce qualities of art at statistically equivalent rates. This is both impractical and undermines the identity-based nature of many art practices. Second, consumers, critics, dealers, and artists alike could be encouraged to reevaluate hierarchies of worth associated with artist categories and features. Specifically, categories and features of art predominantly associated with the work of women, even if they are rare or produced by a small subset

of artists, could be given equal worth to categories and features predominantly associated with the work of men. Such an approach is a cultural endeavor that may be aided by greater transparency and accountability among institutions in what they choose to exhibit, represent, and sell.

This brings me to a final point. Cultural objects, such as art, are the binding of significance to a material form in a way that it remains malleable (Reeves et al., 2015; Taylor et al., 2019). As the social world around those objects changes, so too will their meaning and their markets (Lena, 2006; Lena & Pachucki, 2013; Lopes, 1992; Peterson, 1992). Even beliefs and norms now buttressed by consensus can be broken down with challenge (Ridgeway & Correll, 2006), allowing hierarchies to be inverted, and boundaries between social groups to be made more (or less) permeable, salient, and even visible (DiMaggio, 1987; Lamont & Molnár, 2002). We have seen this already in evidence that suggests women's art is valued more in societies that exhibit more egalitarian gender norms (Adams et al., 2017). An optimistic look from the present dissertation to the future, therefore, accentuates the fact that the art world can, and invariable does, change. The production of scientific findings, such as those given here, are needed to elucidate inequalities so that change can be for the benefit of all.

The three studies of this dissertation contribute to our understanding of how the social status characteristics of producers can impact product markets. I have shown, via the case of gender in the market for contemporary art, that such characteristics can

become salient enough to influence patterns of production and valuation, with the latter being distributed across valuation of labor and valuation of product characteristics. As long theorized and here empirically documented, cultural hierarchies of society are entwined into markets and should be accounted for in empirical research. In grasping how these hierarchies act as structures and instruments of trade, we can we better understand how social inequalities are perpetuated via markets.



## **Appendix A**

### ***A.1 Art Genome Project vis-à-vis academic taxonomies of art***

In its structure, the Art Genome Project is comparable to academic taxonomies for art, but with several important advantages. Academic taxonomies of art, such as the popular Getty Vocabularies, are structured terminologies implemented by museum cataloguers when organizing art and artifact collections. These taxonomies tend to be comprehensive, but they are also interpreted and deployed differently by each institution. As such, even if a social scientist were to collect data from institutions using these taxonomies, the data would not be uniform. Perhaps more importantly, academic taxonomies are specialized for experts and often entail concepts that would not be known or understood by non-art specialists or most art world participants. Because of this, they are unlikely to provide the foundations—if such foundations exist—for inequality between populations of contemporary artists. Finally, the emphasis of academic taxonomies is on ancient artifacts and classical works of art. As such, they are ill-equipped to handle the concepts and mediums present in the Contemporary art world. This contrasts with the Art Genome Project, which, as noted previously, specializes in Contemporary art's evolution.

### ***A.2 Strengths and Limitations of the Art Genome Project***

Are the artworks on Artsy generally representative of the contemporary art world? While a census of artists and works does not exist for comparison, several details

of the data justify caveats for interpreting the generalizability of this study. First, to have artwork listed on Artsy, an artist must be associated with a gallery or museum. It is these organizations that list artists and works on the platform. Artists are not able to list themselves. The primary justification for this mediated process has to do with the historical structure of the offline art market, and with the aims of Artsy as a marketplace for original art.

Since the mid nineteenth century, the art *market* has been organized by a framework in which artists are associated with galleries, or small businesses that represent them and sell their work to interested parties. Their primary role of a gallery is to support the artists they represent by liaising with the buyers and other institutions. This allows artists, theoretically, to focus singularly on their practice. Services provided by galleries include: publicity, hosting exhibitions of the artist's work in solo shows or in group exhibitions that leverage the reputation of several artists, publishing books of an artist's portfolio, forming and maintaining relationships with collectors, managing sales and finances, applying to and participating in art fairs, and ultimately constructing archives.

While the relationship between artists and galleries has been long-standing, online art platforms have emerged over the past decade as new means by which artists can be discovered and sell their work. Several of these platforms have disrupted the norm of gallery-artist relationships by allowing sales directly from artists. Artsy, by

contrast, has built the gallery-artist relationship into the structure and user interface of their platform, and thus acts as a support instead of a disruption to the traditional art world structure. It also allows the platform to more closely mirror the offline contemporary art world.

In addition to mimicking the offline gallery-artist relationship, the injunction for artists to be associated with an art institution on Artsy is a means to specify the type of work listed on the site; namely, original work by professional artists. Galleries function as gatekeepers in the contemporary art world, in that they select which artists and works should be promoted to the public. With few exceptions, galleries only represent original work by professional artists, and exclude reproductions or hobbyist creation. Detriments to this approach include a replication of offline art market biases to online contexts. This should be considered when interpreting findings of the present manuscript. For example, while over fifty percent of art graduates are female, only thirty percent of artists represented by galleries are female. To the degree that women choosing and chosen to be represented by a gallery produce or price art with different characteristics than those who do not, the Artsy platform and Art Genome Project may be biased. A similar caveat applies to early-stage, minority, and other professional but independent artists who are excluded from this study by the requirement of gallery representation to be on Artsy.

Similarly, while Artsy is free for non-profit art institutions, such as museums, most galleries and art fairs are required to pay an annual fee to list their artists and artworks on the site. The exact fees are not in the public domain. The author can only certify that, as it is in Artsy's interest to have a generally representative sample of all galleries on the site, this fee is within a reasonable range for the median art gallery. Nevertheless, in a market that is often uncertain and volatile, cost management is a top concern among gallerist and may prevent them from subscribing to services such as Artsy. From the opposite perspective, Artsy may provide a useful investment for such businesses, as it could lengthen the distance between the physical location of the gallery and they average buyer, and shorten the average time to make a sale after first contact with potential buyers.

Alternative sources of art could augment the present data. For example, one could use the many application programming interfaces (APIs) available for scholars to access art data—including the Google Art Project and APIs associated with individual museums. Beyond being cumbersome to collect data from so many sources, however, these APIs are generally limited. First, they only account for works and artists specific to the institution owning the API and as such are biased toward large, elite institutions with the resources needed to support data dissemination. Second, because museums are more likely to have an API than galleries, these sources are significantly biased toward works produced prior to 1945 and by prominent artists. In the contemporary art world,

by contrast, it is galleries and non-elite artists that comprise the majority of work. Finally, the data available through APIs do not include economic value data for artworks. In point, while the Art Genome Project has limitations and missing data of its own, the unique strengths that it provides produce a novel opportunity for the scholarly investigation of the contemporary art market.

Lastly, aside from the question of which artists and artworks are on Artsy, it is also reasonable to question the bias and completeness of the Art Genome Project. As discussed in the main body of the text, the Art Genome Project is a taxonomy of over 1,000 art-relevant features that are applied to artworks by a collaborative team of engineers and art historians. Importantly for this study, and unlike any other artistic taxonomy presently in existence, the Art Genome Project specializes in the qualities of the contemporary art market—or the market for art produced today or in recent decades. This is a genuinely unprecedentedly opportunity for research, but it's dependence on human coding and the rapid evolution of the art world today renders the data incomplete as new features are added but not easily back-populated to artworks that have been previously coded.

## Appendix B

### *B.1 Selecting artworks for analysis*

The original Artsy data from 2018 included 99,089 unique “artist” cases (including solo and collaborative artists groups), 1,859,637 unique “artwork” cases (including artworks, design and architecture objects, and artifacts), and 8,699 unique “partner” cases (including galleries, museums, and other art institutions).

These original data were filtered to produce the final set for random forest and regression analysis. The following exclusion criteria were applied: (1) artworks created by the collaboration of two or more artists,<sup>8</sup> (2) artworks created by an artist without a name, (3) artworks created by an artist without a gallery association, (4) artworks created prior to 1945 or after 2018<sup>9</sup>, (5) artworks that had not been coded by the Art Genome Project taxonomy, (6) artworks without a value for the “Art” feature in the Art Genome Project taxonomy (the absence of this feature indicates that the object is either an artifact or a design or architecture object), (7) artworks without a non-zero value for at least one “Contemporary” Art Genome Project feature, (8) artworks with a score above 50 for the following features: Film Video, Digital Art and Video, Performance Art, Performing Arts, and (9) artworks with fewer than five non-zero features—where

---

<sup>8</sup> Collaborations were first identified by removing “artist” cases with more than one listed birth year. The remaining set of artists were then hand-coded to identify collaborations.

<sup>9</sup> Artworks without a listed year of creation were included, provided they had a value for at least one “contemporary” feature in the Art Genome Project taxonomy.

features refers to the set of 637 feature used as candidates in Random Forest classification (see Section C1). The largest decrease in sample size results from filter 5, as only 406,988 of the 1,859,637 artworks on Artsy have been coded by the team of art historians and engineers as of 2018.

After applying these filters, the final dataset used for random forest classifier includes: 18,624 solo artists, 255,887 artworks, and 3,384 partner galleries. The final dataset used for regression further limits this data to include only artworks produced after 1999 that are not installations, and that have a gallery listing price greater than 10 but less than maximum outlier of 999,999,999 USD: 102,908 artworks, 11,253 artists, 1,998 galleries.

## ***B.2 Imputation of missing gender data***

Galleries provided gender information for 72% (13,353) of artists in the random forest classifier, and 72% (8,099) of artists in the subset of this data included in the regression. Of 5,474 (28%) artists with missing gender information, gender was first imputed using census records from the United States, United Kingdom, and Great Britain. Specifically, if less than 40% of individuals with a given name were male in the census records, then artists with that given name were identified as female; if more than 70% of individuals with a given name were male in the census records, then artists with that given name were identified as male. Using this method, 4,673 cases of artist gender were imputed. This approach to imputation was tested on the artists records for which

gallery-provided gender information was available, and the predictive accuracy was 95.5%.

For artists still unidentified, I used Google to search art criticism texts and the gallery or artist website—if the artist was described using male pronouns, then the artist was coded as male, if female pronouns were used, then the artist was coded as female. 891 cases of artist gender were imputed.

*Gallery-provided gender:* 13,353 (classifier) 8,099 (regression)

*Census-imputed gender:* 4,673 (classifier) 2,729 (regression)

*Hand-coded gender:* 891 (classifier) 425 (regression)

The only other variable missing information is the number of events (exhibitions and fairs) that an artist participated in. This variable was constructed using Artsy's list of fairs and exhibitions, which is not a census of events for galleries or artists. As such, some galleries and some artists have no events listed on the site, but this should not be interpreted as their having not participated in a fair or exhibition for the timeframe of this study. It may be that these entities chose not to fill or backfill event information in their profiles. Nevertheless, a value of zero was imputed for all galleries and artists with no listed events in the Artsy dataset, as the measure is a proxy for art world involvement.



## Appendix C

### ***C.1 Feature selection and modification for random forest classification***

In the original Art Genome Project taxonomy, there are 1,329 art-relevant features. For the random classifier, features were excluded if they met any of the following exclusion criteria: (1) features that specified gender explicitly (e.g. 'Female master'), (2) features that specified geography, (3) features that specified time period, (4) features that related specifically to art fairs, (5) features that were duplicates. The latter occurred when there were two versions of a feature—a version manually coded by art historians and a version automatically coded by the presence of other features. In this case, missing values in the automatically applied version were filled by values from the manually coded version, and then only the former was kept for analyses. Prior to classifier training, the remaining 634 features were analyzed and removed if, (6) they contained near non-zero variance, or (7) were highly correlated (Pearson correlation coefficient  $> .8$ ), or (8) were represented in the art of fewer than 30 artists. The result is the set of 594 features used in random forest classification.

Justification for removing infrequent features is that they act as noise unrepresentative of population distributions. I remove features appearing in fewer than 30 artist's work, because 30 is the first quartile of the frequency distribution for features by artists. Features related to geography and time period were removed because of their correlation with the location and time period of the artist, which may provide more

information on gender distribution of artists in those geographies and time periods than the gender distribution of artists producing works with qualities associated with those geographies and time periods.

Finally, features were binarized to 0 or 1 from the 0 to 100 scale assigned by the art historian coders. This allows for a score to be present if the artwork displays an artistic feature to any degree, and absent otherwise. The primary motivation for this was to reduce the influence of manual coder variability, and because contact with coders indicated that scores were only applied to a work when substantively present (i.e. very low scores are rare).<sup>10</sup>

## ***C.2 Classifier robustness checks***

Robustness checks for the primary random forest classification method were conducted on the following subsets of data, and with the following methods. Instances 1 through 9 are also presented in Figure 2 of the main manuscript. In each case, 10-fold cross validation and a 70-30 ratio of training to test sets were implemented. All analyses were carried out in R version 1.2.5019, and in most instances using the Caret package version 6.0-85.

- (1) Random Forest classification used in primary analysis
- (2) Random Forest classification using continuous 0 to 100 feature scores: in the primary analyses, feature values were binarized from their original scale of 0-100. This continuous scale is the original form of the Art Genome Project data, with values applied by the team of art historians to indicate how well the feature

---

<sup>10</sup> See sensitivity analysis without binarization in Online Supplement, Section C.2.

is represented in the work. Sensitivity analysis including these continuous values was performed to ensure that results of simplifying features to binarized values of being present or absent from a work did not dramatically alter classificatory accuracy. Results are not substantively altered.

(3) Random Forest classification using only artworks also used in the regression analyses of Section 3.6: this subset analysis was performed to ensure that findings are robust when limiting classification to only artworks that have a listing price and were produced after 1999. Results are not substantively altered.

(4-8) Random Forest classification for each primary medium: primary mediums include paintings, sculpture, photography, textiles, drawing, and other. They are not mutually exclusive, allowing artwork to exist in more than one set. Results are not substantively altered.

(9) Random Forest classification excluding artworks for which the artist's gender was imputed by hand or via census data: this subset analysis was performed to ensure that results are robust when relying only on data in which the artist's gallery identified gender. Results are not substantively altered.

(10) Random Forest with dimensionality reduction: this robustness check relied on Principal Component Analysis to reduce dimensionality of features. Results are not substantively altered.

(11) Lasso: this robustness check relied on LASSO regression, which reduces the influence of uninformative features (noise). Results are not substantively altered.

(12) Logistic regression: this robustness check relied on Logistic regression. Results are not substantively altered.

(13) Boosted tree: this robustness check relied on Boosted tree method of classification, in which iterations of classification aim to predict and reduce the error of the previous. Results are not substantively altered

## Appendix D

### D.1 Mixed effect log linear model of gallery listing price for artwork (Model 2), by medium

Table 11: Mixed effect log linear model of gallery listing price for artwork, by medium (Model 2 in manuscript)

	<i>Painting</i>			<i>Sculpture</i>			<i>Photography</i>			<i>Textile</i>			<i>Drawing</i>		
	Est	Std Err	p	Est	Std Err	p	Est	Std Err	p	Est	Std Err	p	Est	Std Err	Stat
Female	-0.078	0.019	***	-0.104	0.038	***	-0.073	0.037	*	0.546	3.500		0.052	0.044	
<b>Artwork controls</b>															
Size (sq/cu in.)`	0.543	0.002	***	0.304	0.003	***	0.413	0.003	***	0.280	0.010	***	0.496	0.004	***
Year	0.000	0.001	***	-0.006	0.002	***	-0.010	0.001	***	-0.118	0.014	***	0.007	0.002	***
Price hidden on platform	0.052	0.006	***	0.125	0.017	***	0.111	0.014	***	0.404	4.380		0.013	0.014	
<b>Artist controls</b>															
Gallery associations`	0.336	0.020	***	0.382	0.042	***	0.251	0.041	***	2.250	2.910		0.366	0.049	***
Events`	0.075	0.015	***	0.086	0.029	***	0.111	0.030	***	-0.751	2.470		0.035	0.036	
Artworks listed`	-0.141	0.013	***	-0.153	0.027	***	-0.083	0.024	***	1.050	2.400		-0.116	0.033	***
Followers on platform`	0.190	0.011	***	0.175	0.025	***	0.108	0.020	***	-0.783	2.360		0.136	0.028	***
Intercept	3.09	1.77		17.92	4.83	***	25.37	2.76	***	269.75	15893.08		-8.32	3.98	*
Between artist variance		0.42			0.52			0.39			1.54			0.39	
Within artist variance		0.08			0.12			0.06			1334.98			0.06	
BIC		91183.42			40659.81			29164.42			33018.11			23636.24	
N artists (% female)		6,803 (35%)			2,778 (34%)			2,092 (39%)			503 (52%)			1667 (37%)	
N artworks (% female)		59,627 (35%)			15,250 (34%)			14,707 (37%)			2,331 (60%)			10,387 (39%)	
<i>Effective sample size*</i>		35,559			8,420			7,903			1,470			5,847	

\* Models include both covariate balancing propensity score weights and dummy variables for each gallery representing at least one work in the sample

` logged (elasticity) interpretation: 1% change in variable =  $\beta\%$  change in price.

## **Appendix D.2 Robustness check of mixed effect log linear model of gallery listing price for artwork (Model 4)**

A robustness check of Model 4 of this dissertation that includes the addition of statistical controls for whether artwork addresses the themes of *Feminism* or *Gender and Sexual identity*, but that does not subset the data to artwork with non-imputed artist gender information results in a change of statistical significance. Specifically, the interaction effect between observed artist gender and the gender-typicality of the artwork diminishes, leaving no statistically significant difference between women and men in the effect of producing gender-typical features of art. That is, while women's art is still found to be valued less than men's across all values of gender-typicality, and while artworks with more female-typical qualities are still found to be valued less than artworks with more male-typical qualities, these results imply that no statistically significant evidence that the gender gap in listing prices widens or changes with the gender-typicality of features in an artwork.

This deviation from both the primary results of this manuscript and from prior literature may indicate a real and empirical difference in the functioning of the art market, compared to markets for other, less symbolic goods. However, it also may be that this discrepancy is a nuanced artifact of measurement. The only two prior studies to address gender-congruity in product markets operationalized gender typicality at the product level according to how "male" or "female" the products were judged to be by survey participants (Tak et al., 2019), or according to whether women or men

predominantly consume the products (Kricheli-Katz & Regev, 2016). By contrast, the present study measured gender typicality according to whether the features of products are predominantly made by women or men. This approach aligns with many studies of devaluation theory in the labor market, which operationalize the gender typicality of occupations according to the gender of the laborers who perform them. However, it is not attuned to test whether women and men are differentially impacted for producing feminine and masculine qualities of art. It only tests whether they are rewarded for producing art with qualities that are more typical of artists of their gender. Future work should address this discrepancy and advance our understanding of whether and how the value of art made by women and men is altered by their depiction of categories and features that align with expectations of their gender.

## ***Appendix E. Future work on the network structure of the contemporary art market***

A common approach to the study of markets that was not taken in this dissertation is via networks. Such an approach would organize and analyze artists, galleries, and artistic features according to structures that have specifiable relationships and positions according to what these entities do (or do not) have in common. The Artsy data used in this manuscript provide a unique opportunity to investigate the market for contemporary art in this way. Such an approach has been found in previous research to reveal orders of reputation and power, cognitive frames of understanding, and important conduits of information that can shape stability and opportunities in a market. There are several lines of inquiry that could be approached in studying the art market as a network. I abstract two here.


First, the network position of galleries tied by their co-representation of artists, or artists tied by their co-representation by a gallery could provide a sense of the organizational field that was not deeply engaged with in the primary analyses of this dissertation. Even without direct contact, artists' co-membership in a gallery may shape their ability to access material and informational resources that shape artistic practice, audience and institutional reception, or patterns of valuation. Preliminary regression analyses of this organizational structure using the Artsy data and artwork listing price as an outcome (available upon request) suggest that women artists tied to a higher proportion of other women artists in the gallery network are benefited from less dense

networks. That is, women appear to have higher listing prices if the artists they are tied to are not also tied to one another by gallery association. If confirmed, this effect may reflect flows of private information in the art world. Given the highly secretive nature of gallery listing prices, such information flow may be more necessary for women who otherwise exist in information enclaves of lower price setting norms. Alternatively, such open networks may provide women with access to a larger and more diverse audience of consumers, which may increase demand and thereby valuation for their work.

Findings on similar topics have been found by Lutter (2015) in their study of the film industry, and by Yang et al (2019) in their study of post-graduate school leadership placement. Specific to the art world, Banerjee and Ingram (2018) studied the relationship between network structure and art historical fame, as measured by an artist's mentions in Google's n-gram text corpus. The authors find that—even beyond expert and computational measures of creativity—artists working between 1910–1925 were more likely to become famous if they occupied a brokerage (local density) rather than a closure position in this text network. Similarly, in their study of artists from 1981 to 1992, Giuffre (1999) finds that artists in loosely knit social networks receive more critical attention than either artists who have had a long history of membership in tightly knit cliques or those with a history of sporadic connections.



Figure 7 depicts 3,384 galleries representing contemporary artists in the Artsy dataset. Vertices are colored by modularity class and ties are trimmed to only galleries that share two or more artists. Figure 8 depicts the reverse projection of this network, with 6,507 artists tied by their co-representation in a gallery. Vertices are colored by the gender of the artist and sized by the proportion of female artists in their ego network.



**Figure 7. Network representation of art galleries tied by mutual representation of contemporary artists**



**Figure 8. Network representation of contemporary artists tied by mutual representation by at least one gallery**

**Table 12. Artist network variables derived from the network of artists tied by mutual gallery representation**

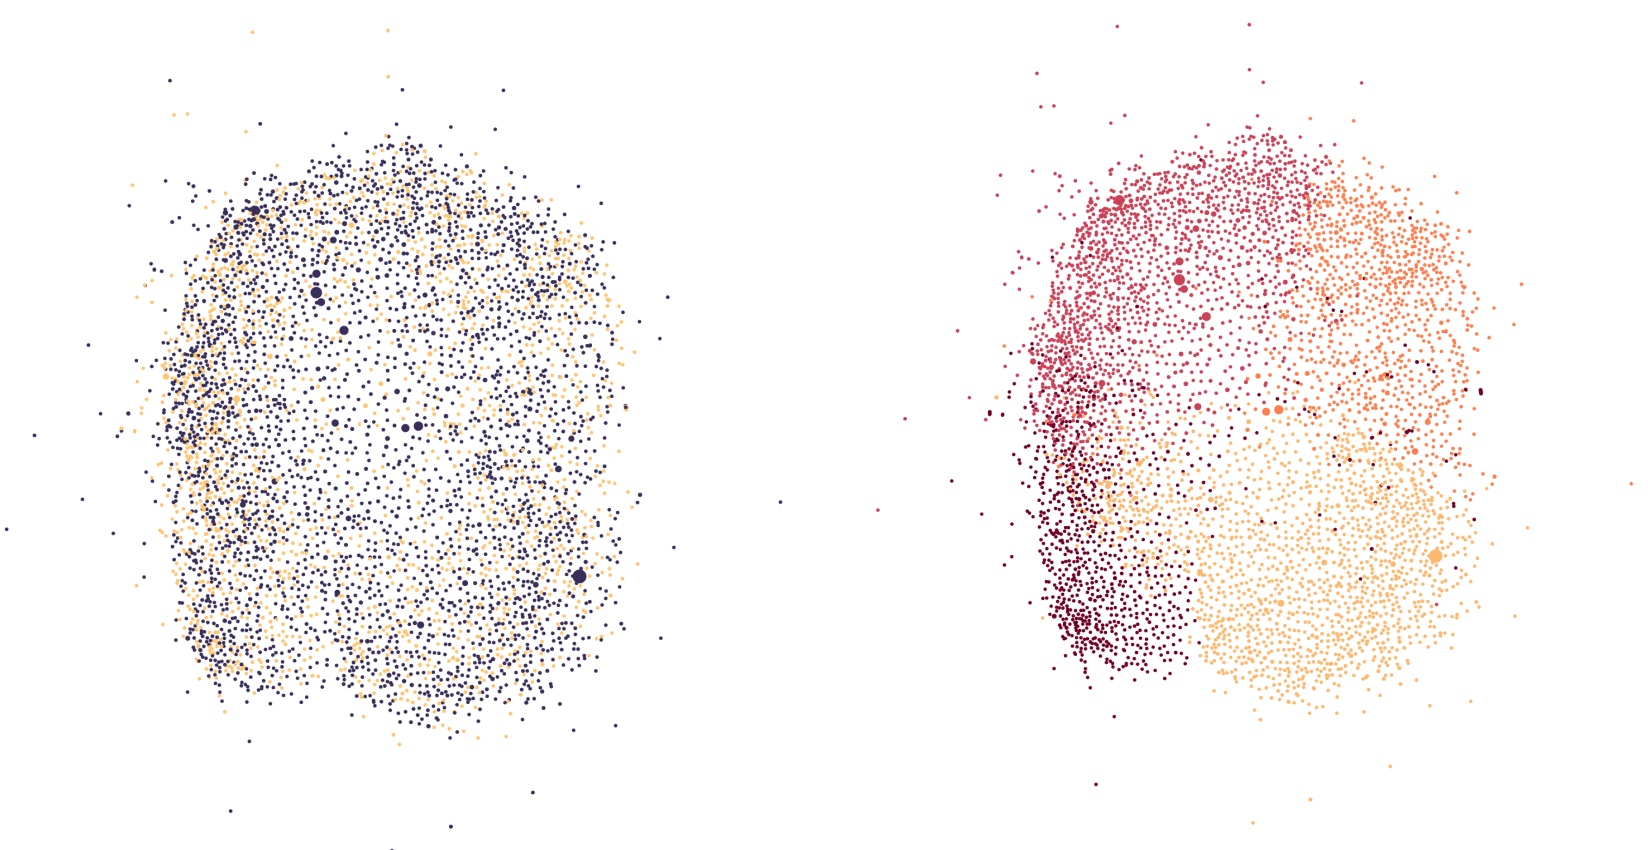
		Overall	Female	Male
<b>Primary network degree</b>	<i>Mean</i>	24.6	22.2	25.9
	<i>Median</i>	16.0	16.0	17.0
	<i>Sd</i>	27.6	22.7	29.8
	<i>Min</i>	1.0	1.0	1.0
	<i>Max</i>	328.0	235.0	328.0
<b>Primary network % female</b>	<i>Mean</i>	30.0	40.0	30.0
	<i>Median</i>	30.0	40.0	30.0
	<i>Sd</i>	20.0	20.0	20.0
	<i>Min</i>	0.0	0.0	0.0
	<i>Max</i>	100.0	100.0	100.0
<b>Secondary network degree</b>	<i>Mean</i>	347.4	293.3	375.9
	<i>Median</i>	147.0	125.0	165.0
	<i>Sd</i>	489.0	409.1	524.1
	<i>Min</i>	0.0	0.0	0.0
	<i>Max</i>	2,992.0	2,825.0	2,992.0
<b>Secondary network % female</b>	<i>Mean</i>	30.0	30.0	30.0
	<i>Median</i>	30.0	30.0	30.0
	<i>Sd</i>	10.0	10.0	10.0
	<i>Min</i>	0.0	0.0	0.0
	<i>Max</i>	100.0	100.0	100.0
<b>Secondary network openness (1-density)</b>	<i>Mean</i>	0.8	0.7	0.8
	<i>Median</i>	0.9	0.9	0.9
	<i>Sd</i>	0.3	0.3	0.3
	<i>Min</i>	0.0	0.0	0.0
	<i>Max</i>	1.0	1.0	1.0

n = 6,507 artists (female = 2,243, male = 4,264)

As an alternative or in addition to networks that tie artists via their gallery representations, artists could be tied via their mutual use of artistic features, or artistic features could be tied via their mutual use by artists. This would provide a sense of the cultural network in which these entities exist. By addressing patterns in the Artsy artistic

feature data via such network relationships, new insights may be gained. For example, contact to other artists using a diverse set of creative features in their art, or positioning between two distinction genres of artistic production (M. A. Pachucki & Breiger, 2010), may alter an artist's ability to innovate or gain rewards for innovation. These effects may differ by gender, given prior evidence of bias against women in assessments of creativity (Proudfoot et al., 2015). This is but one possible line of inquiry that could be gained by structuring artist and artistic feature data as a network of relationships.

Figure 9 depicts contemporary artists, tied by their mutual use of artistic features. For parsimony of visualization, only a random sample 5,000 artists is represented, and the ties used for layout are trimmed to include only features coded at 50 or higher on the scale from 0 to 100 by art historians in the Art Genome Project, and only weights of ten or more—meaning that two artists are tied if they commonly use ten or more features in their art. The graph to the left colors vertices by the gender of the artist, while the graph to the right colors vertices by an estimation of modularity score.



**Figure 9. Network representation of contemporary artists tied by the features represented in their artwork  
(random n = 5,000)**

## References

- Abbott, A. (2010). *Chaos of Disciplines*. University of Chicago Press.
- Adams, R. B., Kraussl, R., Navone, M. A., & Verwijmeren, P. (2017). *Is Gender in the Eye of the Beholder? Identifying Cultural Attitudes with Art Auction Prices* (SSRN Scholarly Paper ID 3083500). Social Science Research Network. <https://papers.ssrn.com/abstract=3083500>
- Akerlof, G. A. (1970). The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500. <https://doi.org/10.2307/1879431>
- Albrecht, J., Björklund, A., & Vroman, S. (2003). Is There a Glass Ceiling in Sweden? *Journal of Labor Economics*, 21(1), 145–177. <https://doi.org/10.1086/344126>
- Alksnis, C., Desmarais, S., & Curtis, J. (2008). Workforce Segregation and the Gender Wage Gap: Is “Women’s” Work Valued as Highly as “Men’s”?1. *Journal of Applied Social Psychology*, 38(6), 1416–1441. <https://doi.org/10.1111/j.1559-1816.2008.00354.x>
- Anderson, J. (1991). The Adaptive Nature Of Human Categorization. *Psychological Review*, 98, 409–429. <https://doi.org/10.1037/0033-295X.98.3.409>
- Askin, N., & Mauskopf, M. (2017). What Makes Popular Culture Popular? Product Features and Optimal Differentiation in Music. *American Sociological Review*, 82(5), 910–944. <https://doi.org/10.1177/0003122417728662>
- Aspers, P. (2008). Knowledge and valuation in markets. *Theory and Society*, 38(2), 111. <https://doi.org/10.1007/s11186-008-9078-9>
- Banerjee, M., & Ingram, P. L. (2018). *Fame as an Illusion of Creativity: Evidence from the Pioneers of Abstract Art* (SSRN Scholarly Paper ID 3258318). Social Science Research Network. <https://papers.ssrn.com/abstract=3258318>
- Baron, J. N., & Newman, A. E. (1990). For What It’s Worth: Organizations, Occupations, and the Value of Work Done by Women and Nonwhites. *American Sociological Review*, 55(2), 155–175. <https://doi.org/10.2307/2095624>
- Baum, J. A. C., & Oliver, C. (1991). Institutional linkages and organizational mortality. *Administrative Science Quarterly*, 36(2), 187-. General Reference Center GOLD.
- Baunach, D. M. (2002). Trends in Occupational Sex Segregation and Inequality, 1950 to 1990. *Social Science Research*, 31(1), 77–98. <https://doi.org/10.1006/ssre.2001.0719>
- Becker, G. S. (1957). *The Economics of Discrimination*. University of Chicago Press.

- Becker, G. S. (1964). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. University of Chicago Press.
- Becker, H. S. (1974). Art As Collective Action. *American Sociological Review*, 39(6), 767–776. <https://doi.org/10.2307/2094151>
- Becker, H. S. (1982). *Art Worlds* (25th Anniversary edition, Updated and Expanded edition 2008). University of California Press.
- Beckert, J. (1996). What Is Sociological about Economic Sociology? Uncertainty and the Embeddedness of Economic Action. *Theory and Society*, 25(6), 803–840. JSTOR.
- Beckert, J. (2009). The social order of markets. *Theory and Society*, 38(3), 245–269. <https://doi.org/10.1007/s11186-008-9082-0>
- Beckert, J. (2019). Markets from meaning: Quality uncertainty and the intersubjective construction of value. *Cambridge Journal of Economics*, bez035. <https://doi.org/10.1093/cje/bez035>
- Beckert, J., & Rössel, J. (2013). The Price of Art. *European Societies*, 15(2), 178–195. <https://doi.org/10.1080/14616696.2013.767923>
- Benjamin, B. A., & Podolny, J. M. (1999). Status, Quality, and Social Order in the California Wine Industry. *Administrative Science Quarterly*, 44(3), 563–589. <https://doi.org/10.2307/2666962>
- Berger, J., Cohen, B. P., & Zelditch, M. (1972). Status Characteristics and Social Interaction. *American Sociological Review*, 37(3), 241–255. JSTOR. <https://doi.org/10.2307/2093465>
- Berger, J., & Fişek, M. H. (2006). Diffuse Status Characteristics and the Spread of Status Value: A Formal Theory. *American Journal of Sociology*, 111(4), 1038–1079. <https://doi.org/10.1086/498633>
- Bergmann, B. R. (1986). *The economic emergence of women*. Basic Books Incorporated.
- Biau, G., & Scornet, E. (2016). A random forest guided tour. *TEST*, 25(2), 197–227. <https://doi.org/10.1007/s11749-016-0481-7>
- Blackburn, R. M., Browne, J., Brooks, B., & Jarman, J. (2002). Explaining gender segregation. *The British Journal of Sociology*, 53(4), 513–536. <https://doi.org/10.1080/0007131022000021461>
- Blau, F. D., Brummund, P., & Liu, A. Y.-H. (2013). Trends in Occupational Segregation by Gender 1970–2009: Adjusting for the Impact of Changes in the Occupational Coding System. *Demography*, 50(2), 471–492. <https://doi.org/10.1007/s13524-012-0151-7>



- Blau, F. D., & Kahn, L. M. (2007). The Gender Pay Gap: Have Women Gone as Far as They Can? *Academy of Management Perspectives*, 21(1), 7–23.
- Blau, F. D., & Kahn, L. M. (2016). *The Gender Wage Gap: Extent, Trends, and Explanations* (IZA DP No. 9656; Discussion Paper Series). Institute for the Study of Labor. <http://ftp.iza.org/dp9656.pdf>
- Blau, P. M. (1964). *Exchange and Power in Social Life*. Transaction Publishers.
- Bocart, F., Gertsberg, M., & Pownall, R. A. J. (2018). *Glass Ceilings in the Art Market* (SSRN Scholarly Paper ID 3079017). Social Science Research Network. <https://papers.ssrn.com/abstract=3079017>
- Boltanski, L., & Thévenot, L. (2006). *On Justification: Economies of Worth*. Princeton University Press.
- Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., & Kalai, A. (2016). Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. *ArXiv:1607.06520 [Cs, Stat]*. <http://arxiv.org/abs/1607.06520>
- Bose, C. E., & Rossi, P. H. (1983). Gender and Jobs: Prestige Standings of Occupations as Affected by Gender. *American Sociological Review*, 48(3), 316–330. <https://doi.org/10.2307/2095225>
- Bourdieu, P. (1979). *Distinction: A Social Critique of the Judgement of Taste* (1 edition). Routledge.
- Bourdieu, P. (1983). The field of cultural production, or: The economic world reversed— ScienceDirect. *Poetics*, 12(4–5), 311–356.
- Bourdieu, P. (1984). The Market of Symbolic Goods\*. In *Pierre Bourdieu The Field of Cultural Production: Essays on Art and Literature*. Columbia University Press.
- Bourdieu, P. (1993). *The Field of Cultural Production* (R. Johnson, Ed.; 1st edition). Columbia University Press.
- Breiman, L. (2001). Random Forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Brescoll, V. L., & Uhlmann, E. L. (2008). Can an angry woman get ahead? Status conferral, gender, and expression of emotion in the workplace. *Psychological Science*, 19(3), 268–275. <https://doi.org/10.1111/j.1467-9280.2008.02079.x>
- Brodersen, K. H., Ong, C. S., Stephan, K. E., & Buhmann, J. M. (2010). The Balanced Accuracy and Its Posterior Distribution. *2010 20th International Conference on Pattern Recognition*, 3121–3124. <https://doi.org/10.1109/ICPR.2010.764>

- Brynin, M., & Perales, F. (2016). Gender Wage Inequality: The De-gendering of the Occupational Structure. *European Sociological Review*, 32(1), 162–174. <https://doi.org/10.1093/esr/jcv092>
- Bryson, B. (1996). “Anything But Heavy Metal”: Symbolic Exclusion and Musical Dislikes. *American Sociological Review*, 61(5), 884–899. <https://doi.org/10.2307/2096459>
- Bryson, B. (1997). What about the univores? Musical dislikes and group-based identity construction among Americans with low levels of education. *Poetics*, 25(2), 141–156. [https://doi.org/10.1016/S0304-422X\(97\)00008-9](https://doi.org/10.1016/S0304-422X(97)00008-9)
- Budig, M. J., & England, P. (2001). The Wage Penalty for Motherhood. *American Sociological Review*, 66(2), 204–225. <https://doi.org/10.2307/2657415>
- Buolamwini, J., & Gebru, T. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *Conference on Fairness, Accountability and Transparency*, 77–91. <http://proceedings.mlr.press/v81/buolamwini18a.html>
- Cameron, L., Goetzmann, W. N., & Nozari, M. (2017). *Art and Gender: Market Bias or Selection Bias?* (SSRN Scholarly Paper ID 3025923). Social Science Research Network. <https://papers.ssrn.com/abstract=3025923>
- Caves, R. E. (2000). *Creative Industries: Contracts Between Art and Commerce*. Harvard University Press.
- Cerulo, K. A. (1988). Analyzing cultural products: A new method of measurement. *Social Science Research*, 17(4), 317–352. [https://doi.org/10.1016/0049-089X\(88\)90007-5](https://doi.org/10.1016/0049-089X(88)90007-5)
- Christin, A. (2012). Gender and highbrow cultural participation in the United States. *Poetics*, 40(5), 423–443. <https://doi.org/10.1016/j.poetic.2012.07.003>
- Coate, B., & Fry, T. R. L. (2012). *Better off Dead? Prices Realised for Australian Paintings Sold at Auction*. <https://www.semanticscholar.org/paper/Better-off-Dead-Prices-Realised-for-Australian-Sold-Coate-Fry/869200f51135dbe7092dc2abc8f9d427f0c9b371>
- Cohen, P. N., & Huffman, M. L. (2007). Working for the Woman? Female Managers and the Gender Wage Gap. *American Sociological Review*, 72(5), 681–704. <https://doi.org/10.1177/000312240707200502>
- Cohen, P. N., Huffman, M. L., & Knauer, S. (2009). Stalled Progress?: Gender Segregation and Wage Inequality Among Managers, 1980-2000. *Work and Occupations*, 36(4), 318–342. <https://doi.org/10.1177/0730888409347582>
- Correll, S. J. (2001). Gender and the Career Choice Process: The Role of Biased Self-Assessments. *American Journal of Sociology*, 106(6), 1691–1730. <https://doi.org/10.1086/321299>

- Correll, S. J. (2004). Constraints into Preferences: Gender, Status, and Emerging Career Aspirations. *American Sociological Review*, 69(1), 93–113. <https://doi.org/10.1177/000312240406900106>
- Correll, S. J., & Ridgeway, C. L. (2006). Expectation States Theory. In *Handbook of Social Psychology*. Springer USA.
- Correll, S. J., Ridgeway, C. L., Zuckerman, E. W., Jank, S., Jordan-Bloch, S., & Nakagawa, S. (2017). It's the Conventional Thought That Counts: How Third-Order Inference Produces Status Advantage. *American Sociological Review*, 82(2), 297–327. <https://doi.org/10.1177/0003122417691503>
- Cowen, T. (1996). Why Women Succeed, and Fail, in the Arts. *Journal of Cultural Economics*, 20(2), 93–113.
- Crino, M. D., White, M. C., & DeSanctis, G. L. (1983). Female participation rates and the occupational prestige of the professions: Are they inversely related? *Journal of Vocational Behavior*, 22(2), 243–255. [https://doi.org/10.1016/0001-8791\(83\)90030-1](https://doi.org/10.1016/0001-8791(83)90030-1)
- Danto, A. (1964). The artworld. *The Journal of Philosophy*, 61(19), 571–584.
- de la Fuente, E. (2007). The 'New Sociology of Art': Putting Art Back into Social Science Approaches to the Arts. *Cultural Sociology*, 1(3), 409–425. <https://doi.org/10.1177/1749975507084601>
- Dey, S. (2015). Impact of remittances on poverty at origin: A study on rural households in India using Covariate Balancing Propensity Score Matching. *Migration and Development*, 4(2), 185–199. <https://doi.org/10.1080/21632324.2014.979022>
- DiMaggio, P. (1987). Classification in Art. *American Sociological Review*, 52(4), 440–455. <https://doi.org/10.2307/2095290>
- DiMaggio, P. (1997). Culture and Cognition. *Annual Review of Sociology*, 23, 263–287.
- DiMaggio, P., & Mukhtar, T. (2004). Arts participation as cultural capital in the United States, 1982–2002: Signs of decline? *Poetics*, 32(2), 169–194. <https://doi.org/10.1016/j.poetic.2004.02.005>
- DiMaggio, P., & Mukhtar, T. (2018). Arts participation as cultural capital in the United States, 1982–2002: Signs of decline? *Poetics*, 32(2), 169–194.
- Dinovitzer, R., Reichman, N., & Sterling, J. (2009). The Differential Valuation of Women's Work: A New Look at the Gender Gap in Lawyers' Incomes. *Social Forces*, 88(2), 819–864. <https://doi.org/10.1353/sof.0.0260>
- Du Bois, W. E. B. (Ed.). (1899). *The Philadelphia Negro: A Social Study*.

- Durkheim, É. (1915). *The elementary forms of the religious life, a study in religious sociology*. G. Allen & Unwin; Macmillan.  
<https://search.library.wisc.edu/catalog/999467601502121/cite>
- Eagly, A. H., & Karau, S. J. (2002). Role congruity theory of prejudice toward female leaders. *Psychological Review*, *109*(3), 573–598.
- Edwards, S. (2004). The Economics of Latin American Art: Creativity Patterns and Rates of Return. *Economía*, *4*(2), 1–35. <https://doi.org/10.1353/eco.2004.0015>
- Elias, N., & Scotson, J. L. (1965). *The Established and the Outsiders: A Sociological Enquiry Into Community Problems*. Cass & Company.
- England, P. (1979). Women and Occupational Prestige: A Case of Vacuous Sex Equality. *Signs*, *5*(2), 252–265.
- England, P. (1992). *Comparable Worth: Theories and Evidence*. Transaction Publishers.
- England, P. (2005). Gender Inequality in Labor Markets: The Role of Motherhood and Segregation. *Social Politics: International Studies in Gender, State & Society*, *12*(2), 264–288. <https://doi.org/10.1093/sp/jxi014>
- England, P., Allison, P., & Wu, Y. (2007). Does bad pay cause occupations to feminize, Does feminization reduce pay, and How can we tell with longitudinal data? *Social Science Research*, *36*(3), 1237–1256. <https://doi.org/10.1016/j.ssresearch.2006.08.003>
- England, P., Farkas, G., Kilbourne, B. S., & Dou, T. (1988). Explaining Occupational Sex Segregation and Wages: Findings from a Model with Fixed Effects. *American Sociological Review*, *53*(4), 544–558. <https://doi.org/10.2307/2095848>
- England, P., Herbert, M. S., Kilbourne, B. S., Reid, L. L., & Megdal, L. M. (1994). The Gendered Valuation of Occupations and Skills: Earnings in 1980 Census Occupations. *Social Forces*, *73*(1), 65–100. <https://doi.org/10.2307/2579918>
- Espeland, W. N., & Sauder, M. (2007). Rankings and Reactivity: How Public Measures Recreate Social Worlds. *American Journal of Sociology*, *113*(1), 1–40.  
<https://doi.org/10.1086/517897>
- Fine, G. A. (1996). Reputational Entrepreneurs and the Memory of Incompetence: Melting Supporters, Partisan Warriors, and Images of President Harding. *American Journal of Sociology*, *101*(5), 1159–1193.
- Fiske, S. T., & Glick, P. (1996). The Ambivalent Sexism Inventory: Differentiating hostile and benevolent sexism. *Journal of Personality and Social Psychology*, *70*(3), 491–512.
- Fligstein, N., & Dauter, L. (2007). The Sociology of Markets. *Annual Review of Sociology*, *33*(1), 105–128. <https://doi.org/10.1146/annurev.soc.33.040406.131736>

- Flores, A. R., & Barclay, S. (2016). Backlash, Consensus, Legitimacy, or Polarization: The Effect of Same-Sex Marriage Policy on Mass Attitudes. *Political Research Quarterly*, 69(1), 43–56. <https://doi.org/10.1177/1065912915621175>
- Flores, A. R., Haider-Markel, D. P., Lewis, D. C., Miller, P. R., Tadlock, B. L., & Taylor, J. K. (2018). Challenged Expectations: Mere Exposure Effects on Attitudes About Transgender People and Rights. *Political Psychology*, 39(1), 197–216. <https://doi.org/10.1111/pops.12402>
- Folleco, A., Khoshgoftaar, T. M., Van Hulse, J., & Bullard, L. (2008). Software quality modeling: The impact of class noise on the random forest classifier. *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*, 3853–3859. <https://doi.org/10.1109/CEC.2008.4631321>
- Fombrun, C., & Shanley, M. (1990). What's in a Name? Reputation Building and Corporate Strategy. *Academy of Management Journal; Briarcliff Manor*, 33(2), 233.
- Foschi, M. (2000). Double Standards for Competence: Theory and Research. *Annual Review of Sociology*, 26(1), 21–42. <https://doi.org/10.1146/annurev.soc.26.1.21>
- Fourcade, M. (2007). Theories of Markets and Theories of Society. *American Behavioral Scientist*, 50(8), 1015–1034. <https://doi.org/10.1177/0002764207299351>
- Fourcade, M. (2011). Cents and Sensibility: Economic Valuation and the Nature of “Nature.” *American Journal of Sociology*, 116(6), 1721–1777. <https://doi.org/10.1086/659640>
- Fourcade, M., & Healy, K. (2017). Categories All the Way Down. *Historical Social Research / Historische Sozialforschung*, 42(1 (159)), 286–296.
- Frerichs, S. (2019). Bounded sociality: Behavioural economists' truncated understanding of the social and its implications for politics. *Journal of Economic Methodology*, 26(3), 243–258. <https://doi.org/10.1080/1350178X.2019.1625217>
- Frith, S. (1996). Music and Identity. In S. Hall & P. Du. Gay (Eds.), *Questions of Cultural Identity* (pp. 108–125). Sage.
- Frow, J. (1995). *Cultural Studies and Cultural Value* (1 edition). Clarendon Press.
- Frow, J. (2014). *Genre* (2 edition). Routledge.
- Gaytán, M. S. (2008). From Sombreros to Sincronizadas: Authenticity, Ethnicity, and the Mexican Restaurant Industry. *Journal of Contemporary Ethnography*, 37(3), 314–341. <https://doi.org/10.1177/0891241607309621>
- Giuffre, K. (1999). Sandpiles of Opportunity: Success in the Art World. *Social Forces*, 77(3), 815–832. <https://doi.org/10.2307/3005962>

- Godart, F. C., & Galunic, C. (2019). Explaining the Popularity of Cultural Elements: Networks, Culture, and the Structural Embeddedness of High Fashion Trends. *Organization Science*, 30(1), 151–168. <https://doi.org/10.1287/orsc.2018.1234>
- Goldberg, A. (2011). Mapping Shared Understandings Using Relational Class Analysis: The Case of the Cultural Omnivore Reexamined. *American Journal of Sociology*, 116(5), 1397–1436. <https://doi.org/10.1086/657976>
- Goldin, C., & Rouse, C. (2000). Orchestrating Impartiality: The Impact of “Blind” Auditions on Female Musicians. *American Economic Review*, 90(4), 715–741. <https://doi.org/10.1257/aer.90.4.715>
- Gorman, E. H., & Kmec, J. A. (2009). Hierarchical Rank and Women’s Organizational Mobility: Glass Ceilings in Corporate Law Firms. *American Journal of Sociology*, 114(5), 1428–1474. <https://doi.org/10.1086/595950>
- Gould, R. V. (2002). The Origins of Status Hierarchies: A Formal Theory and Empirical Test. *American Journal of Sociology*, 107(5), 1143–1178. <https://doi.org/10.1086/341744>
- Greenfeld, L. (1988). Professional Ideologies and Patterns of “Gatekeeping”: Evaluation and Judgment within Two Art Worlds. *Social Forces*, 66(4), 903–925. <https://doi.org/10.2307/2579428>
- Guy, M. E., & Newman, M. A. (2004). Women’s Jobs, Men’s Jobs: Sex Segregation and Emotional Labor. *Public Administration Review*, 64(3), 289–298. <https://doi.org/10.1111/j.1540-6210.2004.00373.x>
- Hall, E. L. (1993). WAITERING/WAITRESSING: Engendering the Work of Table Servers. *Gender & Society*, 7(3), 329–346. <https://doi.org/10.1177/089124393007003002>
- Halle, D. (1996). *Inside Culture: Art and Class in the American Home*. University of Chicago Press.
- Heilman, M. E. (2001). Description and Prescription: How Gender Stereotypes Prevent Women’s Ascent Up the Organizational Ladder. *Journal of Social Issues*, 57(4), 657–674. <https://doi.org/10.1111/0022-4537.00234>
- Heilman, M. E. (2012). Gender stereotypes and workplace bias. *Research in Organizational Behavior*, 32, 113–135. <https://doi.org/10.1016/j.riob.2012.11.003>
- Heilman, M. E., Wallen, A. S., Fuchs, D., & Tamkins, M. M. (2004). Penalties for success: Reactions to women who succeed at male gender-typed tasks. *The Journal of Applied Psychology*, 89(3), 416–427. <https://doi.org/10.1037/0021-9010.89.3.416>
- Hewitt, M. P. (2015). Self-Efficacy, Self-Evaluation, and Music Performance of Secondary-Level Band Students. *Journal of Research in Music Education*, 63(3), 298–313. <https://doi.org/10.1177/0022429415595611>

- Holt, F. (2007). *Genre in Popular Music*. University of Chicago Press.
- Homans, G. C. (1961). *Social Behavior: Its Elementary Forms*. Harcourt, Brace and World, Inc.
- Hsu, G. (2006). Jacks of All Trades and Masters of None: Audiences' Reactions to Spanning Genres in Feature Film Production. *Administrative Science Quarterly*, 51(3), 420–450. <https://doi.org/10.2189/asqu.51.3.420>
- Hsu, G., Hannan, M. T., & Koçak, Ö. (2009). Multiple Category Memberships in Markets: An Integrative Theory and Two Empirical Tests. *American Sociological Review*, 74(1), 150–169. <https://doi.org/10.1177/000312240907400108>
- Huang, C. (2013). Gender differences in academic self-efficacy: A meta-analysis. *European Journal of Psychology of Education*, 28(1), 1–35. <https://doi.org/10.1007/s10212-011-0097-y>
- Huffman, M. L., & Velasco, S. C. (1997). When More Is Less. Sex Composition, Organizations, and Earnings in U.S. Firms. *Work and Occupations: An International Sociological Journal*, 24(2), 214–244.
- Hughes, E. C. (1945). Dilemmas and Contradictions of Status. *American Journal of Sociology*, 50(5), 353–359. <https://doi.org/10.1086/219652>
- Hume, D. (1757). *Of the Standard of Taste: Selected Essays*. Oxford University Press.
- Imai, K., & Ratkovic, M. (2014). Covariate balancing propensity score. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 76(1), 243–263. JSTOR.
- Jensen, M. (2008). The Use of Relational Discrimination to Manage Market Entry: When Do Social Status and Structural Holes Work Against You? *Academy of Management Journal*, 51(4), 723–743. <https://doi.org/10.5465/amr.2008.33665259>
- Johnson, G., & Solon, G. (1986). Estimates of the Direct Effects of Comparable Worth Policy. *The American Economic Review*, 76(5), 1117–1125.
- Jones, C., Maoret, M., Massa, F. G., & Svejenova, S. (2011). Rebels with a Cause: Formation, Contestation, and Expansion of the De Novo Category “Modern Architecture,” 1870–1975. *Organization Science*, 23(6), 1523–1545. <https://doi.org/10.1287/orsc.1110.0701>
- Karpik, L. (2010). Valuing the Unique: The Economics of Singularities. In *Economics Books*. Princeton University Press. <https://ideas.repec.org/b/pup/pbooks/9215.html>
- Katz-Gerro, T., & Sullivan, O. (2002). Leisure, Tastes and Gender in Britain: Changes from the 1960s to the 1990s\*. *Undefined*. <https://www.semanticscholar.org/paper/Leisure%2C-Tastes-and-Gender-in-Britain%3A-Changes-from-Katz-Gerro-Sullivan/55af363c6f465c8c6b6bbbed6866a811ba237bf0>

- Kilbourne, B. S., England, P., Farkas, G., Beron, K., & Weir, D. (1994). Returns to Skill, Compensating Differentials, and Gender Bias: Effects of Occupational Characteristics on the Wages of White Women and Men. *American Journal of Sociology*, *100*(3), 689–719.
- Kim, J. W., & King, B. G. (2014). Seeing Stars: Matthew Effects and Status Bias in Major League Baseball Umpiring. *Management Science*, *60*(11), 2619–2644. <https://doi.org/10.1287/mnsc.2014.1967>
- King, G., & Nielsen, R. (2019). Why Propensity Scores Should Not Be Used for Matching. *Political Analysis*, *27*(4), 435–454. <https://doi.org/10.1017/pan.2019.11>
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., & Mullainathan, S. (2018). Human Decisions and Machine Predictions\*. *The Quarterly Journal of Economics*, *133*(1), 237–293. <https://doi.org/10.1093/qje/qjx032>
- Ko, I., Kotrba, L., & Roebuck, A. (2015). Leaders as Males?: The Role of Industry Gender Composition. *Sex Roles*, *72*(7–8), 294–307. <https://doi.org/10.1007/s11199-015-0462-4>
- Koçak, Ö. (2003). *Social orders of exchange: Effects and origins of social order in exchange markets*. Stanford University.
- Koenker, R., & Bassett, G. (1978). Regression Quantiles. *Econometrica*, *46*(1), 33–50. <https://doi.org/10.2307/1913643>
- Kovács, B., & Hannan, M. T. (2015). Conceptual Spaces and the Consequences of Category Spanning. *Sociological Science*, *2*, 252–286. <https://doi.org/10.15195/v2.a13>
- Kricheli-Katz, T., & Regev, T. (2016). How many cents on the dollar? Women and men in product markets. *Science Advances*, *2*(2), e1500599. <https://doi.org/10.1126/sciadv.1500599>
- Kricheli-Katz, T., Regev, T., & Correll, S. (2019). Why Are Women Penalized in Product Markets? *Socius*, *5*, 2378023119861024. <https://doi.org/10.1177/2378023119861024>
- Ku, M. C. (2011). When Does Gender Matter?: Gender Differences in Specialty Choice Among Physicians. *Work and Occupations*, *38*(2), 221–262. <https://doi.org/10.1177/0730888410392319>
- Lamont, M. (2012). Toward a Comparative Sociology of Valuation and Evaluation. *Annual Review of Sociology*, *38*(1), 201–221. <https://doi.org/10.1146/annurev-soc-070308-120022>
- Lamont, M., & Molnár, V. (2002). The Study of Boundaries in the Social Sciences. *Annual Review of Sociology*, *28*, 167–195.
- Lang, G. E., & Lang, K. (1988). Recognition and Renown: The Survival of Artistic Reputation. *American Journal of Sociology*, *94*(1), 79–109.



- Lena, J. C. (2006). Social Context and Musical Content of Rap Music, 1979–1995. *Social Forces*, 85(1), 479–495. <https://doi.org/10.1353/sof.2006.0131>
- Lena, J. C. (2012). *Banding Together: How Communities Create Genres in Popular Music*. Princeton University Press.
- Lena, J. C., & Pachucki, M. C. (2013). The sincerest form of flattery: Innovation, repetition, and status in an art movement. *Poetics*, 41(3), 236–264. <https://doi.org/10.1016/j.poetic.2013.02.002>
- Lena, J. C., & Peterson, R. A. (2008). Classification as Culture: Types and Trajectories of Music Genres. *American Sociological Review*, 73(5), 697–718. <https://doi.org/10.1177/000312240807300501>
- Levanon, A., England, P., & Allison, P. (2009). Occupational Feminization and Pay: Assessing Causal Dynamics Using 1950–2000 U.S. Census Data. *Social Forces*, 88(2), 865–891. <https://doi.org/10.1353/sof.0.0264>
- Lizardo, O., & Skiles, S. (2016). Cultural Objects as Prisms: Perceived Audience Composition of Musical Genres as a Resource for Symbolic Exclusion. *Socius*, 2, 2378023116641695. <https://doi.org/10.1177/2378023116641695>
- Lopes, P. D. (1992). Innovation and Diversity in the Popular Music Industry, 1969 to 1990. *American Sociological Review*, 57(1), 56–71. <https://doi.org/10.2307/2096144>
- Lundberg, S., & Rose, E. (2000). Parenthood and the earnings of married men and women. *Labour Economics*, 7(6), 689–710. [https://doi.org/10.1016/S0927-5371\(00\)00020-8](https://doi.org/10.1016/S0927-5371(00)00020-8)
- Lutter, M. (2015). Do Women Suffer from Network Closure? The Moderating Effect of Social Capital on Gender Inequality in a Project-Based Labor Market, 1929 to 2010. *American Sociological Review*, 80(2), 329–358. <https://doi.org/10.1177/0003122414568788>
- Lynn, F. B., Podolny, J. M., & Tao, L. (2009). A Sociological (De)Construction of the Relationship between Status and Quality. *American Journal of Sociology*, 115(3), 755–804. <https://doi.org/10.1086/603537>
- Macpherson, D. A., & Hirsch, B. T. (1995). Wages and Gender Composition: Why do Women's Jobs Pay Less? *Journal of Labor Economics*, 13(3), 426–471. <https://doi.org/10.1086/298381>
- Magnusson, C. (2009). Gender, Occupational Prestige, and Wages: A Test of Devaluation Theory. *European Sociological Review*, 25(1), 87–101.
- Magnusson, C. (2010). Why Is There a Gender Wage Gap According to Occupational Prestige?: An Analysis of the Gender Wage Gap by Occupational Prestige and Family Obligations in Sweden. *Acta Sociologica*, 53(2), 99–117. <https://doi.org/10.1177/0001699310365627>

- Marx, K. (1867). *Das Kapital* (F. Engles, Ed.). Regnery Publishin.
- Mason, W., Vaughan, J. W., & Wallach, H. (2014). Computational social science and social computing. *Machine Learning*, 95(3), 257–260. <https://doi.org/10.1007/s10994-013-5426-8>
- McAndrew, C. (Ed.). (2019). *The Art Market 2019* (Art Basel & UBS Report). Arts Economics. <https://www.artbasel.com/amr2019>
- McAndrew, C. (Ed.). (2020). *The Art Market 2020* (Art Basel & UBS Report). Arts Economics. <https://www.artbasel.com/amr2020>
- McAndrew, C. (Ed.). (2021). *The Art Market 2021* (Art Basel & UBS Report). Arts Economics. <https://www.artbasel.com/amr2020>
- Menger, P.-M. (1999). Artistic Labor Markets and Careers. *Annual Review of Sociology*, 25(1), 541–574. <https://doi.org/10.1146/annurev.soc.25.1.541>
- Merton, R. K. (1988). The Matthew Effect in Science, II: Cumulative Advantage and the Symbolism of Intellectual Property. *Isis*, 79(4), 606–623. <https://doi.org/10.1086/354848>
- Moss-Racusin, C. A., Dovidio, J. F., Brescoll, V. L., Graham, M. J., & Handelsman, J. (2012). Science faculty’s subtle gender biases favor male students. *Proceedings of the National Academy of Sciences*, 109(41), 16474–16479. <https://doi.org/10.1073/pnas.1211286109>
- Mullainathan, S., & Obermeyer, Z. (2017). Does Machine Learning Automate Moral Hazard and Error? *American Economic Review*, 107(5), 476–480. <https://doi.org/10.1257/aer.p20171084>
- Mullainathan, S., & Spiess, J. (2017). Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives*, 31(2), 87–106. <https://doi.org/10.1257/jep.31.2.87>
- National Endowment for the Arts. (2019). *Artists and Other Cultural Workers: A Statistical Portrait*. [https://www.arts.gov/sites/default/files/Artists\\_and\\_Other\\_Cultural\\_Workers.pdf](https://www.arts.gov/sites/default/files/Artists_and_Other_Cultural_Workers.pdf)
- Nochlin, L. (1971). Why Have There Been No Great Women Artists? In V. Gornick & B. Moran (Eds.), *Woman in Sexist Society: Studies in Power and Powerlessness*. Basic Books.
- Nunes, J. C., & Ordanini, A. (2014). I like the way it sounds: The influence of instrumentation on a pop song’s place in the charts. *Musicae Scientiae*, 18(4), 392–409. <https://doi.org/10.1177/1029864914548528>
- Ochsenfeld, F. (2014). Why Do Women’s Fields of Study Pay Less? A Test of Devaluation, Human Capital, and Gender Role Theory. *European Sociological Review*, 30(4), 536–548. <https://doi.org/10.1093/esr/jcu060>

- Ortega, O. D., Øvergård, K. I., & Henden, V. (2015). Women Are Better Leaders Than They Think: Gender Differences in the Self-Assessment of Leadership Skills in the Maritime Industry. In M. Kitada, E. Williams, & L. L. Froholdt (Eds.), *Maritime Women: Global Leadership* (pp. 201–216). Springer. [https://doi.org/10.1007/978-3-662-45385-8\\_15](https://doi.org/10.1007/978-3-662-45385-8_15)
- Pachucki, M. A., & Breiger, R. L. (2010). Cultural Holes: Beyond Relationality in Social Networks and Culture. *Annual Review of Sociology*, *36*(1), 205–224. <https://doi.org/10.1146/annurev.soc.012809.102615>
- Pachucki, M. C. (2012). Classifying quality: Cognition, interaction, and status appraisal of art museums. *Poetics*, *40*(1), 67–90. <https://doi.org/10.1016/j.poetic.2011.11.001>
- Padgett, J. F., & Ansell, C. K. (1993). Robust Action and the Rise of the Medici, 1400-1434. *American Journal of Sociology*, *98*(6), 1259–1319.
- Park, D. Y., & Podolny, J. M. (2000). The competitive dynamics of status and niche width: US investment banking, 1920-1949. *Industrial and Corporate Change*, *9*(3), 377–414. <https://doi.org/10.1093/icc/9.3.377>
- Peterson, R. A. (1992). Understanding audience segmentation: From elite and mass to omnivore and univore. *Poetics*, *21*(4), 243–258. [https://doi.org/10.1016/0304-422X\(92\)90008-Q](https://doi.org/10.1016/0304-422X(92)90008-Q)
- Peterson, R. A., & Anand, N. (2004). The Production of Culture Perspective. *Annual Review of Sociology*, *30*(1), 311–334. <https://doi.org/10.1146/annurev.soc.30.012703.110557>
- Peterson, R. A., & Kern, R. M. (1996). Changing Highbrow Taste: From Snob to Omnivore. *American Sociological Review*, *61*(5), 900–907. <https://doi.org/10.2307/2096460>
- Petrusich, A. (2021, March 15). Genre is Disappearing. What Comes Next? *New Yorker*.
- Phillips, D. J. (2001). The Promotion Paradox: Organizational Mortality and Employee Promotion Chances in Silicon Valley Law Firms, 1946–1996. *American Journal of Sociology*, *106*(4), 1058–1098. <https://doi.org/10.1086/320293>
- Podolny, J. M. (1993a). A Status-Based Model of Market Competition. *American Journal of Sociology*, *98*(4), 829–872.
- Podolny, J. M. (1993b). A Status-Based Model of Market Competition. *American Journal of Sociology*, *98*(4), 829–872.
- Podolny, J. M. (1994). Market uncertainty and the social character of economic exchange. *Administrative Science Quarterly*, *39*(3), 458–483.
- Podolny, J. M., & Phillips, D. J. (1996). The Dynamics of Organizational Status. *Industrial and Corporate Change*, *5*(2), 453–471. <https://doi.org/10.1093/icc/5.2.453>

- Pontikes, E., & Hannan, M. (2014). An Ecology of Social Categories. *Sociological Science, 1*, 311–343. <https://doi.org/10.15195/v1.a20>
- Posner, R. A. (1993). *Cardozo: A Study in Reputation*. University of Chicago Press.
- Pownall, R. (2017). *TEFAF 2017 Global Art Market Report*. The European Fine Art Foundation. <https://www.tefaf.com/about/art-market-report>
- Proudfoot, D., Kay, A. C., & Koval, C. Z. (2015). A Gender Bias in the Attribution of Creativity: Archival and Experimental Evidence for the Perceived Association Between Masculinity and Creative Thinking. *Psychological Science, 26*(11), 1751–1761. <https://doi.org/10.1177/0956797615598739>
- Reeves, A., Gilbert, E., & Holman, D. (2015). Class dis-identification, cultural stereotypes, and music preferences: Experimental evidence from the UK. *Poetics, 50*, 44–61. <https://doi.org/10.1016/j.poetic.2015.01.002>
- Reskin, B. F. (1993). Sex Segregation in the Workplace. *Annual Review of Sociology, 19*, 241–270.
- Reskin, B. F., & Roos, P. A. (1990). *Job Queues, Gender Queues: Explaining Women's Inroads Into Male Occupations*. Temple University Press.
- Riach, P. A., & Rich, J. (2002). Field Experiments of Discrimination in the Market Place. *The Economic Journal, 112*(483), F480–F518.
- Ridgeway, C. L. (2001). Gender, Status, and Leadership. *Journal of Social Issues, 57*(4), 637–655. <https://doi.org/10.1111/0022-4537.00233>
- Ridgeway, C. L., & Correll, S. J. (2004). Unpacking the Gender System: A Theoretical Perspective on Gender Beliefs and Social Relations. *Gender & Society, 18*(4), 510–531. <https://doi.org/10.1177/0891243204265269>
- Ridgeway, C. L., & Correll, S. J. (2006). Consensus and the Creation of Status Beliefs. *Social Forces, 85*(1), 431–453. <https://doi.org/10.1353/sof.2006.0139>
- Ridgeway, C. L., & Smith-Lovin, L. (1999). The Gender System and Interaction. *Annual Review of Sociology, 25*(1), 191–216. <https://doi.org/10.1146/annurev.soc.25.1.191>
- Rosen, S. (1981). The Economics of Superstars. *The American Economic Review, 71*(5), 845–858.
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika, 70*(1), 41–55. <https://doi.org/10.1093/biomet/70.1.41>

- Rosenfeld, R. A. (1980). Race and Sex Differences in Career Dynamics. *American Sociological Review*, 45(4), 583–609. <https://doi.org/10.2307/2095010>
- Rudman, L. A., & Phelan, J. E. (2008). Backlash effects for disconfirming gender stereotypes in organizations. *Research in Organizational Behavior*, 28, 61–79. <https://doi.org/10.1016/j.riob.2008.04.003>
- Rule, A., & Bearman, P. (2015). Networks and Culture. In L. Hanquinet & M. Savage (Eds.), *Routledge International Handbook of the Sociology of Art and Culture* (1st ed.). Routledge.
- Saito, T., & Rehmsmeier, M. (2015). The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets. *PLOS ONE*, 10(3), e0118432. <https://doi.org/10.1371/journal.pone.0118432>
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science*, 311(5762), 854–856. <https://doi.org/10.1126/science.1121066>
- Salganik, M. J., Lundberg, I., Kindel, A. T., Ahearn, C. E., Al-Ghoneim, K., Almaatouq, A., Altschul, D. M., Brand, J. E., Carnegie, N. B., Compton, R. J., Datta, D., Davidson, T., Filippova, A., Gilroy, C., Goode, B. J., Jahani, E., Kashyap, R., Kirchner, A., McKay, S., ... McLanahan, S. (2020). Measuring the predictability of life outcomes with a scientific mass collaboration. *Proceedings of the National Academy of Sciences*, 117(15), 8398–8403. <https://doi.org/10.1073/pnas.1915006117>
- Salganik, M. J., Lundberg, I., Kindel, A. T., & McLanahan, S. (2019). Introduction to the Special Collection on the Fragile Families Challenge. *Socius*, 5, 2378023119871580. <https://doi.org/10.1177/2378023119871580>
- Sauder, M. (2006). Third parties and status position: How the characteristics of status systems matter. *Theory and Society; Dordrecht*, 35(3), 299–321. <http://dx.doi.org.proxy.lib.duke.edu/10.1007/s11186-006-9005-x>
- Sauder, M., Lynn, F., & Podolny, J. M. (2012). Status: Insights from Organizational Sociology. *Annual Review of Sociology*, 38(1), 267–283. <https://doi.org/10.1146/annurev-soc-071811-145503>
- Simcoe, T. S., & Waguespack, D. M. (2010). Status, Quality, and Attention: What's in a (Missing) Name? *Management Science*, 57(2), 274–290. <https://doi.org/10.1287/mnsc.1100.1270>
- Simmel, G. (1957). Fashion. *American Journal of Sociology*, 62(6), 541–558.
- Sonnett, J. (2004). Musical boundaries: Intersections of form and content. *Poetics*, 32(3), 247–264. <https://doi.org/10.1016/j.poetic.2004.05.007>

- Spence, A. M. (1974). *Market signaling: Informational transfer in hiring and related screening processes*. Cambridge : Harvard University Press, 1974.  
<https://find.library.duke.edu/catalog/DUKE000107564>
- Steele, J., James, J. B., & Barnett, R. C. (2002). Learning in a Man's World: Examining the Perceptions of Undergraduate Women in Male-Dominated Academic Areas. *Psychology of Women Quarterly*, 26(1), 46–50. <https://doi.org/10.1111/1471-6402.00042>
- Steinberg, R. (2001). Comparable Worth in Gender Studies. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social and Behavioral Sciences* (Vol. 4, pp. 2293–2297). Pergamon.
- Steinpreis, R. E., Anders, K. A., & Ritzke, D. (1999). The Impact of Gender on the Review of the Curricula Vitae of Job Applicants and Tenure Candidates: A National Empirical Study. *Sex Roles*, 41(7–8), 509–528. <https://doi.org/10.1023/A:1018839203698>
- Stuart, T. E., & Ding, W. W. (2006). When Do Scientists Become Entrepreneurs? The Social Structural Antecedents of Commercial Activity in the Academic Life Sciences. *American Journal of Sociology*, 112(1), 97–144. <https://doi.org/10.1086/502691>
- Stuart, T. E., Hoang, H., & Hybels, R. C. (1999a). Interorganizational Endorsements and the Performance of Entrepreneurial Ventures. *Administrative Science Quarterly*, 44(2), 315–349. <https://doi.org/10.2307/2666998>
- Stuart, T. E., Hoang, H., & Hybels, R. C. (1999b, June 1). *Interorganizational endorsements and the performance of entrepreneurial ventures*. *Administrative Science Quarterly*.  
<http://link.galegroup.com/apps/doc/A55306276/AONE?sid=googlescholar>
- Suchner, R. W. (1979). Sex Ratios and Occupational Prestige: Three Failures to Replicate a Sexist Bias. *Personality and Social Psychology Bulletin*, 5(2), 236–239.  
<https://doi.org/10.1177/014616727900500224>
- Tak, E., Correll, S. J., & Soule, S. A. (2019). *Gender Inequality in Product Markets: When and How Status Beliefs Transfer to Products*. <https://doi.org/10.1093/sf/soy125>
- Tam, T. (1997). Sex Segregation and Occupational Gender Inequality in the United States: Devaluation or Specialized Training? *American Journal of Sociology*, 102(6), 1652–1692. <https://doi.org/10.1086/231129>
- Taylor, M. A., Stoltz, D. S., & McDonnell, T. E. (2019). Binding significance to form: Cultural objects, neural binding, and cultural change. *Poetics*, 73, 1–16.  
<https://doi.org/10.1016/j.poetic.2019.01.005>
- Thornton, P. H. (1999). The Sociology of Entrepreneurship. *Annual Review of Sociology*, 25(1), 19–46. <https://doi.org/10.1146/annurev.soc.25.1.19>

- Tomaskovic-Devey, D. (1993). *Gender and Racial Inequality at Work: The Sources and Consequences of Job Segregation* (1 edition). ILR Press.
- Topaz, C. M., Klingenberg, B., Turek, D., Heggeseth, B., Harris, P. E., Blackwood, J. C., Chavoya, C. O., Nelson, S., & Murphy, K. M. (2018). Diversity of Artists in Major U.S. Museums. *ArXiv:1812.03899 [Stat]*. <http://arxiv.org/abs/1812.03899>
- Torlò, V. J., & Lomi, A. (2017). The Network Dynamics of Status: Assimilation and Selection. *Social Forces*, *96*(1), 389–422. <https://doi.org/10.1093/sf/sox040>
- Touhey, J. C. (1974a). Effects of Additional Men on Prestige and Desirability of Occupations Typically Performed by Women. *Journal of Applied Social Psychology*, *4*(4), 330–335. <https://doi.org/10.1111/j.1559-1816.1974.tb02604.x>
- Touhey, J. C. (1974b). Effects of additional women professionals on ratings of occupational prestige and desirability. *Journal of Personality and Social Psychology*, *29*(1), 86–89. <https://doi.org/10.1037/h0035742>
- Touw, W. G., Bayjanov, J. R., Overmars, L., Backus, L., Boekhorst, J., Wels, M., & van Hijum, S. A. F. T. (2013). Data mining in the Life Sciences with Random Forest: A walk in the park or lost in the jungle? *Briefings in Bioinformatics*, *14*(3), 315–326. <https://doi.org/10.1093/bib/bbs034>
- Townsend, R. B. (2017). Taking Note: How About Those Undergraduate Arts Majors? *National Endowment for the Arts*. <https://www.arts.gov/art-works/2017/taking-note-how-about-those-undergraduate-arts-majors>
- Tversky, A. (1977). Features of Similarity. *Psychology Review*, *84*, 327–352. <https://doi.org/10.1037/0033-295X.84.4.327>
- Valentino, L. (2020). The Segregation Premium: How Gender Shapes the Symbolic Valuation Process of Occupational Prestige Judgments. *Social Forces*, *99*(1), 31–58. <https://doi.org/10.1093/sf/soz145>
- Veblen, T. (1899). *The Theory of the Leisure Class*. Transaction Publishers.
- Velthuis, O. (2003). Symbolic Meanings of Prices: Constructing the Value of Contemporary Art in Amsterdam and New York Galleries. *Theory and Society*, *32*(2), 181–215.
- Velthuis, O. (2005). *Talking Prices: Symbolic Meanings of Prices on the Market for Contemporary Art*. Princeton University Press.
- Wagner, D. G., & Berger, J. (1997). Gender and Interpersonal Task Behaviors: Status Expectation Accounts. *Sociological Perspectives*, *40*(1), 1–32.
- Waldfoegel, J. (1997). The Effect of Children on Women's Wages. *American Sociological Review*, *62*(2), 209–217. <https://doi.org/10.2307/2657300>

- Waldron, A., Miller, D. C., Redding, D., Mooers, A., Kuhn, T. S., Nibbelink, N., Roberts, J. T., Tobias, J. A., & Gittleman, J. L. (2017). Reductions in global biodiversity loss predicted from conservation spending. *Nature*, *551*(7680), 364–367. <https://doi.org/10.1038/nature24295>
- Washington, M., & Zajac, E. J. (2005). Status Evolution and Competition: Theory and Evidence. *The Academy of Management Journal*, *48*(2), 282–296. <https://doi.org/10.2307/20159656>
- Weber, M. (1978). *Economy and Society: An Outline of Interpretive Sociology*. University of California Press.
- Weichselbaumer, D., & Winter-Ebmer, R. (2005). A Meta-Analysis of the International Gender Wage Gap. *Journal of Economic Surveys*, *19*(3), 479–511.
- White, H. C. (1981). Where Do Markets Come From? *American Journal of Sociology*, *87*(3), 517–547.
- White, H. C., & White, C. A. (1993). *Canvases and Careers: Institutional Change in the French Painting World* (1 edition). University Of Chicago Press.
- White, M. C., Crino, M. D., & DeSanctis, G. L. (1981). Ratings of Prestige and Desirability: Effects of Additional Women Entering Selected Business Occupations. *Personality and Social Psychology Bulletin*, *7*(4), 588–592. <https://doi.org/10.1177/014616728174011>
- White, M. D. (2017). Preferences All the Way Down: Questioning the Neoclassical Foundations of Behavioral Economics and Libertarian Paternalism. *Æconomia. History, Methodology, Philosophy*, *7–3*, 353–373. <https://doi.org/10.4000/oeconomia.2646>
- Wijnberg, N. M., & Gemser, G. (2000). Adding Value to Innovation: Impressionism and the Transformation of the Selection System in Visual Arts. *Organization Science*, *11*(3), 323–329. <https://doi.org/10.1287/orsc.11.3.323.12499>
- Wyss, R., Ellis, A. R., Brookhart, M. A., Girman, C. J., Jonsson Funk, M., LoCasale, R., & Stürmer, T. (2014). The Role of Prediction Modeling in Propensity Score Estimation: An Evaluation of Logistic Regression, bCART, and the Covariate-Balancing Propensity Score. *American Journal of Epidemiology*, *180*(6), 645–655. <https://doi.org/10.1093/aje/kwu181>
- Xu, W., & Leffler, A. (1992). Gender and Race Effects on Occupational Prestige, Segregation, and Earnings. *Gender & Society*, *6*(3), 376–392. <https://doi.org/10.1177/089124392006003003>
- Yang, Y., Chawla, N. V., & Uzzi, B. (2019). A network's gender composition and communication pattern predict women's leadership success. *Proceedings of the National Academy of Sciences*, *116*(6), 2033–2038. <https://doi.org/10.1073/pnas.1721438116>



- Yarkoni, T., & Westfall, J. (2017). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science*, 12(6), 1100–1122. <https://doi.org/10.1177/1745691617693393>
- Zelizer, V. A. (1979). *Morals and Markets: The Development of Life Insurance in the United States*. Columbia University Press.
- Zelizer, V. A. (1981). The Price and Value of Children: The Case of Children's Insurance. *American Journal of Sociology*, 86(5), 1036–1056. <https://doi.org/10.1086/227353>
- Zelizer, V. A. (2000). The Purchase of Intimacy. *Law & Social Inquiry*, 25(3), 817–848. <https://doi.org/10.1111/j.1747-4469.2000.tb00162.x>
- Zolberg, V. L. (2015). A cultural sociology of the arts. *Current Sociology*, 63(6), 896–915. <https://doi.org/10.1177/0011392114551652>

## Biography

Taylor Whitten Brown began her PhD at Duke University in 2015, where she specialized in Networks and Computational Social Science. She was the recipient of the James B. Duke Fellowship and the Katherine Goodman Stern Fellowship. Prior to her time at Duke, Taylor earned a Master of Arts degree in Sociology from the University of North Carolina-Chapel Hill in 2015. She also earned a Master of Science degree in Evidence-based Social Intervention from the University of Oxford in 2011, where she was given a mark Distinction on her final thesis, later published as “*Centre-based Day Care for Children Under-5 and their Families in Low- and Middle-income Countries (A Systematic Review)*” in the Cochran Database of Systematic Reviews (2014). Taylor graduated Magna Cum Laude from Brigham Young University in 2009, where she obtained a Bachelor of Arts degree in Humanities-Philosophy, with a minor in Sociology.

Taylor has co-authored several studies, including “*Channeling Hearts and Minds: Advocacy Organizations, Cognitive-Emotional Currents, and Public Conversation on Social Media*” in the American Sociological Review (2017), “*Exposure to Opposing Views can Increase Political Polarization: Evidence from a Large- Scale Field Experiment on Social Media*” in the Proceedings of the National Academies of Science (2018), “*Power, Proximity, and Prejudice: A Tardean Analysis of the Global Diffusion of Google Search Terms*” in the American Journal of Sociology (2019), and “*Overcrowding and COVID-19 Mortality Across U.S. Counties: Are Disparities Growing Over Time?*” in Social Science and Medicine – Public Health (2021). She has also completed two PhD internships on the Core Data Science team at Facebook, and completed an eighteen month appointment as Science Assistant at the National Science Foundation in Washington DC.