Quantitative Sentiment Analysis: Applications in Finance and Contemporary Art

D I S S E R T A T I O N
of the University of St. Gallen,
School of Management,
Economics, Law, Social Science,
International Affairs and Computer Science,
to obtain the title of
Doctor of Philosophy in Economics and Finance

submitted by

Jan Serwart

from

Zuwil (St.Gallen)

Approved on the application of

Prof. Francesco Audrino, Ph.D.

and

Prof. Dr. Jens Beckert

Dissertation no. 5416

Difo-Druck GmbH, Untersiemau 2024

Quantitative Sentiment Analysis: Applications in Finance and Contemporary Art

D I S S E R T A T I O N
of the University of St. Gallen,
School of Management,
Economics, Law, Social Science,
International Affairs and Computer Science,
to obtain the title of
Doctor of Philosophy in Economics and Finance

submitted by

Jan Serwart

from

Zuzwil (St.Gallen)

Approved on the application of

Prof. Francesco Audrino, PhD

and

Prof. Dr. Jens Beckert
Prof. Dr. Ulrich Matter
Prof. Guido Cozzi, PhD

Dissertation no. 5416

Difo-Druck GmbH, Untersiemau 2024

The University of St.Gallen, School of Management, Economics, Law, Social Sciences, International Affairs and Computer Science, hereby consents to the printing of the present dissertation, without hereby expressing any opinion on the views herein expressed.

St. Gallen, November 15, 2023

The President:

Prof. Dr. Bernhard Ehrenzeller

"This is water."

David Foster Wallace

I thank my supervisor Francesco Audrino, who endowed me with the necessary freedom in various dimensions to craft this dissertation. Furthermore, I would like to thank my parents, my beloved ones and my friends for their constant support. Finally, I would like to thank Florian Paul Koenig, who inspired me to write two papers on the contemporary art market.

Summary

The first paper investigates the predictive power of macro sentiment indicators for the U.S. Treasury yield curve. We use sentiment data constructed by RavenPack and we also create a novel macro sentiment dataset by exploiting text information gathered from the social media platform Twitter. We apply a sentiment classification approach to determine a sentiment score per tweet. Furthermore, we obtain monthly yield curve shape factors by estimating the dynamic Nelson-Siegel model. Using a vector autoregressive and a Markov-switching autoregressive model, we predict changes in the shape of the yield curve. We exploit our shape change predictions in a trading strategy that uses a butterfly approach. We find some evidence that models exploiting inflation sentiment are economically useful when trading the curvature of the yield curve.

The second paper introduces a methodology to dynamically quantify the reputation of contemporary artists. The newly generated data reflect how much status-weighted attention various artists receive from a huge network of art market representatives. The paper develops three results. First, art galleries act as gatekeepers for success in the contemporary art market. Second, other subgroups of the contemporary art market like curators, art advisers, media, collectors, museums and others follow the lead of art galleries in terms of attention granted to an artist. Third, the value of artworks at auctions is impacted by the amount of status-weighted attention the artist has received from art galleries, art advisors and media, with the influence of galleries being the strongest, followed by art advisors and media.

The third paper builds upon the results of the second paper. Specifically, the paper derives that past dynamics of collector artist attention is predictive of future price appreciations for the artworks of an artist. In a trading application, it is shown that choosing artists based on their past collector attention dynamics generates better financial outcomes compared to a chance selection of artists. The suggested approach also reduces the risk of financial loss when investing in contemporary art by avoiding purchase of artworks by artists who have poor recent histories of collector attention dynamics compared to their peers.

Zusammenfassung

Im ersten Aufsatz wird die Vorhersagekraft von Makro-Stimmungsindikatoren für die Renditekurve von US-Staatsanleihen untersucht. Wir verwenden Stimmungsdaten, die von RavenPack erstellt wurden, und erstellen außerdem einen neuartigen Makro-Stimmungsdatensatz, indem wir Textinformationen nutzen, die von der Social-Media-Plattform Twitter gesammelt wurden. Wir wenden einen Sentiment-Klassifizierungsansatz an, um einen Sentiment-Score pro Tweet zu ermitteln. Darüber hinaus ermitteln wir monatliche Renditekurvenformfaktoren durch Schätzung des dynamischen Nelson-Siegel-Modells. Mithilfe eines autoregressiven Vektor- und eines autoregressiven Markov-Switching-Modells sagen wir Veränderungen in der Form der Renditekurve voraus. Wir nutzen unsere Vorhersagen der Formveränderungen in einer Handelsstrategie, die einen "Butterfly" Ansatz verwendet. Wir finden Hinweise darauf, dass Modelle, die die Inflationsstimmungsdaten ausnutzen, beim Handel mit der Krümmung der Renditekurve wirtschaftlich nützlich sind.

Im zweiten Beitrag wird eine Methodik zur dynamischen Quantifizierung des Ansehens zeitgenössischer Künstler vorgestellt. Die neu generierten Daten spiegeln wider, wie viel statusgewichtete Aufmerksamkeit verschiedene Künstler von einem grossen Netzwerk von Kunstmarktvertretern erhalten. Der Aufsatz liefert drei Ergebnisse. Erstens: Kunstgalerien fungieren als Gatekeeper für den Erfolg auf dem zeitgenössischen Kunstmarkt. Zweitens folgen andere Untergruppen des zeitgenössischen

Kunstmarktes wie Kuratoren, Kunstberater, Medien, Sammler, Museen und andere dem Beispiel der Kunstgalerien in Bezug auf die einem Künstler gewährte Aufmerksamkeit. Drittens wird der Wert von Kunstwerken bei Auktionen durch den Umfang der statusgewichteten Aufmerksamkeit beeinflusst, die der Künstler von Kunstgalerien, Kunstberatern und Medien vor der Auktion erhalten hat, wobei der Einfluss der Galerien am stärksten ist, gefolgt von Kunstberatern und Medien.

Das dritte Papier baut auf den Ergebnissen des zweiten Papiers auf. Insbesondere wird gezeigt, dass die vergangene Dynamik der Aufmerksamkeit von Sammlern für Künstler nützlich für die Vorhersage betreffend zukünftigen Preissteigerungen für die Kunstwerke eines Künstlers ist. In einer Handelsanwendung wird gezeigt, dass die Auswahl von Künstlern auf der Grundlage ihrer Aufmerksamkeitsdynamik bei Sammlern zu besseren finanziellen Ergebnissen führt als eine zufällige Auswahl von Künstlern. Der vorgeschlagene Ansatz verringert auch das Risiko finanzieller Verluste bei Investitionen in zeitgenössische Kunst, indem Kunstwerke von Künstler nicht gekauft werden, die in der jüngeren Vergangenheit im Vergleich zu ihren Mitbewerbern, eine schlechtere Aufmerksamkeitsdynamik bei Kunstsammlern und Kunstsammlerinnen hatten.

Contents

Summary **vii**

Zusammenfassung **ix**

1 Yield Curve Trading Strategies Exploiting Sentiment Data **1**

1.1 Introduction 3

1.2 Literature Review 5

1.3 Data 8

1.3.1 Measuring Sentiment 8

1.3.2 Relation to Macroeconomic data 12

1.3.3 Yield Curve Data 14

1.3.4 Monetary policy indicators 16

1.4 Statistical Models 17

1.4.1 Baseline Model 17

1.4.2 Extension to include sentiment variables and monetary policy indicators 20
1.4.3 Markov switching vector autoregressive models . . 21
1.5 Trading Strategy Description 22
1.6 Empirical Results 26
1.6.1 Discussion Term Structure Models 26
1.6.2 Discussion trading results 33
1.7 Conclusion 40

References **42**

Appendix **48**
1.A Attributes and Summary Statistics for Sentiment Variables 48
1.B RavenPack Sentiment vs. Macroeconomic Variables . . . 49
1.C Complete List of Models 51
1.D Performance of Raw Models 52

2 Quantifying Artist Reputation: Disentangling Art-Market Subgroup Credentialing Effects **53**
2.1 Introduction 55
2.2 Literature Review 59
2.3 Credentialing Process Model 65
2.4 Data 67
2.4.1 Mapping the Global Art Community 70
2.5 Methodology 75
2.5.1 Measuring Status-Weighted Artist Attention . . . 75

2.5.2 Models 78
2.6 Results 83
2.6.1 Attention Dynamics 83
2.6.2 Price Levels of Contemporary Art 90
2.6.3 Disentangling Credentialing Effects 96
2.7 Conclusion 102

References **104**

Appendix **110**

2.A Top 100 Gallery and Museum Instagram Usernames . . . 110
2.B Data-Driven Discount Factors 113
2.C Relative Importance Subgroup Representatives 115
2.D Curator Attention and Prices 116
2.E Random Subset Robustness Test 117
2.F Artist Fixed Effects Price Regression 118
2.G Alternative Lagged Price Specifications 119
2.H Additional Instagram Descriptive Statistics 121

3 How to optimally select contemporary artists from an investment perspective **123**

3.1 Introduction 125
3.2 Credentialing Process Model 126
3.3 Methodology 128
3.3.1 Mapping the Global Art Community 128
3.3.2 Measuring Status-Weighted Artist Attention . . . 129

3.3.3 Models 131
3.4 Data 134
3.4.1 In-Sample and Out-of-Sample Data 137
3.5 Results and Application 138
3.6 Conclusion 145

References **147**

Appendix **149**
3.A Top 100 Gallery and Museum Instagram Usernames . . . 149
3.B Return Descriptive Statistics - At Least 1 Painting 152
3.C Return Descriptive Statistics - At Least 4 Paintings . . . 153
3.D Robustness Test 4 Paintings per Year: Regression 154
3.E Robustness Test 4 Paintings per Year: Trading Application 155

Curriculum Vitae **159**

Chapter 1

Yield Curve Trading Strategies Exploiting Sentiment Data

Revise & Resubmit at North American Journal of Economics and Finance

Co-authors: Francesco Audrino

Abstract

This paper builds upon previous research findings that show macro sentiment data-augmented models are better at predicting the yield curve. We extend the dynamic Nelson-Siegel model with macro sentiment data from either Twitter or RavenPack. Vector autogressive (VAR) models and Markov-switching VAR models are used to predict changes in the shape of the yield curve. We build bond butterfly trading strategies that exploit our yield curve shape change predictions. Although the economic returns from our trading strategies based upon models exploiting macro sentiment data do not statistically significantly differ from those which do not rely on it, we find some evidence that models exploiting inflation sentiment are economically useful when trading the curvature of the yield curve.

Keywords: bond butterflies, yield curve, sentiment data

1.1 Introduction

The idea that asset prices are in part driven by investors' psychology has been under academic scrutiny since the early 1990s; see e.g. Daniel et al. (2002) for a literature review. Recent advances in data sciences and the ubiquitous use of social media platforms by financial market participants allowed for the creation of novel market sentiment data. These data complement the survey data on consumers' expectations, which had been available long before the advent of the Internet. The way sentiment data can be used to explain financial market developments has attracted significant scholarly attention in previous years. The potential of sentiment data to explain stock market developments is much better understood compared to fixed income markets. A better understanding of how sentiment data can describe fixed income market developments is warranted given its significance in terms of overall value, connectedness to the real economy and its effects on governments and private entities with interest rate exposure.

This paper extends the existing literature on the applicability of sentiment data in fixed income markets. Audrino and Offner (2022) show that sentiment data statistically significantly increase the predictive power of simple Taylor rule-based interest rate models as well as term structure models of the Nelson-Siegel family and arbitrage-free, short rate-based models. However, thus far no concrete economic application has been tested, that would show that the use of sentiment data also translates into significantly higher economic returns. This paper fills the gap.

We use a term-structure model of the Nelson-Siegel family to estimate yield curve factors. Applying simple vector autoregressive models (VAR) and Markov-switching VAR (MSVAR) models augmented with macroeconomic sentiment data and monetary policy indicators, we forecast the Nelson-Siegel yield curve factors to predict changes in the shape of the yield curve. We exploit these forecasts using a bond butterfly trading strategy.

The empirical results show that the inclusion of sentiment data in models to forecast changes in the shape of the yield curve can have economic benefits. However, concerning bets on the slope of the yield curve, the outperformance of sentiment-augmented models over non-sentiment-augmented models is neither statistically significant nor robust over time. Nevertheless, we find mild evidence that models exploiting inflation sentiment are economically useful when trading the curvature of the yield curve. Inflation sentiment-augmented models consistently rank among the top models for all considered trading periods and deliver economic returns that outperform the benchmark, especially during the Coronavirus pandemic. This paper proceeds as follows: Section 1.2 summarizes related literature, followed by Section 1.3 describing the data. Section 1.4 presents the statistical models, which is followed by Section 1.5 introducing the trading strategy. Finally, Section 1.6 discusses the empirical results. Section 1.7 concludes.

1.2 Literature Review

This paper contributes generally to the term structure literature given impetus by Diebold and Li (2006), who first showed that the time series of the factors from a dynamic Nelson-Siegel model (DNS) contain information that can be exploited for yield curve forecasting purposes using low order vector autoregressions (VAR). Diebold et al. (2006) extended the basic Nelson-Siegel approach to include observable macroeconomic variables. However, the evidence concerning the superior forecasting power of these macroeconomic variable-augmented DNS models remains somewhat mixed (Duffee, 2011). Recently, a series of models was introduced to account for the well-known stylized fact that yields are subject to regime shifts; for an overview, see Filipova et al. (2014). For example, Xiang and Zhu (2013) show that Markov regime-switching DNS models for the U.S. yield curve have superior predictive performance compared to single-state DNS models. Hevia et al. (2014) showed similarly in a no-arbitrage context that allowing for Markov regimes in DNS models can entail superior predictive performance for some parametrizations compared to single regime DNS VAR models. Guidolin and Pedio (2019) drop the no-arbitrage requirements and show that Markov-switching VAR models augmented with monetary policy indicators outperform models without such indicators in terms of the forecast accuracy of U.S. yield curve. So far, however, little research has appeared on whether these models can benefit from the inclusion of macro sentiment data. Modern sentiment data is often constructed using methods from the field of natural language processing (NLP).

The advances in the field of NLP have been significant in past years due to increasing computational power. One stream of NLP research is devoted to the development of methods and models to harness quantitative sentiment data from raw text to be used in empirical analysis. Recent survey articles on the developments in this field include Algaba et al. (2020) and Renault (2019). The progress in the field of NLP has led to an increase in studies researching the potential of sentiment data to predict asset price movements. Most notably, the effectiveness of sentiment data created on the basis of text from social media platforms like Twitter or Stocktwits has gained a lot of attention (Audrino et al. 2020; Bollen et al. 2011; Nofer & Hinz 2015; Oliveira et al. 2013; Renault 2017; Sul et al. 2016; Sun et al. 2016; Tan & Tas 2020).

The sentiment literature pertaining to interest rate developments is less developed. Early work focused on sentiment analysis of central bank communication mediums such as FOMC statements. For example, Lucca and Trebbi (2009) find that the tone of policy communications influences longer-dated treasury yields. Other non-central bank-related sources for sentiment analysis in the realm of interest rate research include the general news. In this regard, RavenPack data have become frequently used. RavenPack marks every news item from a broad universe of news sources with a sentiment value. For example, Erlwein-Sayer (2017) using RavenPack data investigates the relationship between news sentiments and European sovereign yield spreads. Audrino and Offner (2022) show that augmented Nelson-Siegel (NS) models with sentiment data from RavenPack have superior predictive performance when forecasting the U.S. yield curve for different maturities.

Moreover, other studies focus directly on specific macroeconomic data releases and their effects on interest rates. In this respect, Edison (1997) researches the reaction of short- and long-term treasury yields to unemployment and inflation-related news. She finds that positive inflation deviations from expectations increase yields while positive unemployment deviations from expectations lower yields. Goldberg and Leonard (2005) show that the largest moves in yields are driven by U.S. announcements on labor market conditions, real GDP growth, and consumer sentiment.

Concerning yield curve movements, the contribution by Wright (2011) found that the term structure of yields reacts to changes in inflation uncertainty as measured by variation in monetary policy frameworks. Furthermore, Goldberg and Grisse (2013) find evidence that the reaction of the U.S. Treasury yield curve to economic data is not constant over time. Recently, Gotthelf and Uhl (2018) introduced the idea of a novel sentiment factor besides the well-known level-slope-curvature factors in a Nelson-Siegel context. They show that sentiment data improve the prediction accuracy of yields especially at the short end of the curve.

While the literature on forecasting treasury yield curves is well established, few authors have developed concrete trading strategies that exploit the forecasting power of their models. Fabozzi et al. (2005) use the forecasts on changes in the shape parameters of the Nelson-Siegel model in order to build swap butterflies. These are among the most commonly used active trading strategies to exploit predictions concerning the shape of the yield curve. They show that swap butterflies based on their model earn economically significant returns, but they do not statistically benchmark their results against other forecasting models. Recently, Guidolin and Pedio (2019) used a bond butterfly trading strategy to show that the superior statistical performance of their yield curve forecasting models that include monetary policy indicators to predict changes in the shape of the yield curve can translate into higher trading profits. However, the realized returns are very uncertain, and as a result they cannot statistically confirm return differences across models. This paper extends the literature with a concrete trading application based on DNS factor forecasts from MSVAR models augmented with macro sentiment data from diverse sources.

1.3 Data

1.3.1 Measuring Sentiment

We use three different sources to create macroeconomic sentiment data: RavenPack, Twitter, and the Michigan consumer sentiment surveys.

RavenPack News Analytics provides a global macro dataset, of which the

Event Sentiment Score (ESS) is used as sentiment variable. The Event Sentiment Score is calculated by applying a supervised Bayes classifier trained with news articles labelled by financial experts as having either a positive or negative short-term effect on the economy. The algorithm is applied on news articles published by Dow Jones Newswires, Wall Street Journal, Barron's, and Market Watch. The ESS sentiment scores range from 0 (negative) to (100) positive. The RavenPack sentiment scores have been found to contain information exploitable for forecasting purposes for German Bund yields, various European sovereign spreads, and CDS spreads (Erlwein-Sayer, 2017, 2018; Yang et al., 2020).

The macroeconomic variables assumed to have an effect on the yield curve are inflation, unemployment and the Federal Funds Rate (FFR), which we also denote short rate. We build sentiment variables concerning inflation, unemployment and the short-term interest rate. For this choice, we follow Audrino and Offner (2022), who have shown that such sentiment variables have predictive power for the yield curve. For the sake of comparability, we use the same procedure as Audrino and Offner (2022) to construct the sentiment variables. Specifically, for all articles related to one of the macroeconomic variables, we download its ESS scores from RavenPack. In a next step, we aggregate the ESS scores for each macroeconomic variable monthly by taking the mean. Thus, the sentiment score can be stated as:

$$S_h(t) = \frac{1}{N}\sum_{j=1}^{N} SS(k_{j,h})\mathbb{1}_{\{Month(k_{j,h})=t\}}, \qquad \text{for } h = i, \pi, u, \tag{1.1}$$

where SS stands for one ESS sentiment score, $k_{j,h}$ corresponds to one news item about h, where i, π, and u stand for interest rates, inflation and unemployment rate, respectively. N is the number of observed articles in a month, and t denotes the respective month. The applied keywords and summary statistics are given in Appendix 1.A in Table 1.9.

Twitter is a social media network with roughly 400 million active daily users. On the platform, people can share thoughts, ideas, and opinions in the form of short messages currently consisting of currently 280 characters.[1] The data are obtained using an academic API terminal provided by Twitter. For each of the macroeconomic variables inflation, interest rates and unemployment, respectively, we download all tweets that contain specific keywords related to the underlying macroeconomic variables. The keywords and sentiment summary statistics are given in Appendix 1.A in Table 1.10. Due to academic Twitter API limitations, the entirety of keywords used for the RavenPack data set could not be used for the Twitter dataset. However, since at least the eponymous macroeconomic variable keyword was used, we capture the macroeconomic sentiment sufficiently well. The econometric results do not contradict this assumption. Overall, 4.35 million tweets were analyzed. In order to assign a sentiment value to

[1] Before November 7, 2017, the limit was 140 characters.

each tweet, we use a dictionary approach. Dictionaries are a set of words labelled to express a certain sentiment. Specifically, the sentiment of a text is determined by counting the words therein that also appear in the dictionary. In this paper, we use the dictionary provided by Loughran and McDonald (2011), henceforth LM. In an extensive analysis of different approaches to gauging investor sentiment, Ballinari and Behrendt (2021) show that the easy to implement and transparent dictionary approach performs best in an asset pricing application. Moreover, the models estimated using the LM dictionary belong to Hansen model confidence set together with those using the finance-specific dictionary by Renault (2017). For every tweet, we count the number of positive and negative words used in the Twitter post as given by the LM dictionary. The sentiment of a given social media message is calculated as the difference between the share of positive and negative words occuring in the text data. As for the RavenPack dataset, we aggregate the individual tweet sentiment scores for each macroeconomic variable monthly by taking the mean. Hence, the sentiment can be calculated analogously as in equation (1.1).

Finally, we use the *University of Michigan Consumer Sentiment Index* as a third sentiment dataset. This dataset contains information pertaining to consumer expectations for the economy. The sentiment score reflects people's expectations concerning their current financial health, as well as the health of the economy in the short- and long-term. The sentiment scores cannot be disaggregated at the level of our macroeconomic variables but reflect sentiment towards the economy as a whole.

1.3.2 Relation to Macroeconomic data

The potential of sentiment data to improve the forecast accuracy of our analysis depends crucially on the assumption that the extracted sentiment data bear a relationship to the underlying macroeconomic variables that affect interest rates. Figure 1.1 shows the relationship between the macroeconomic variables and their Twitter sentiment counterpart. We observe a statistically significant positive correlation between the change of short-term interest rates and its contemporaneous sentiment counterpart as depicted in Table 1.1. This is in sharp contrast to the results when using RavenPack data as depicted in Appendix 1.B, where we find a negative relationship between those two variables. Audrino and Offner (2022) argue that the correlation is negative for the RavenPack data because "[the RavenPack] sentiment variable is mainly based on the perception of news based directly on interest rates. As declining interest rates boost the performance of fixed income securities, it is possible that investors will rate as positive news which reports falling interest rates" (p. 20). They reason in such a way because other studies such as those by Gotthelf and Uhl (2018), who find a positive relationship, are based on all articles spanning the entire Thomson Reuters News Analytics universe, which are both mostly related to business cycles. Following this logic, the tweets from Twitter on interest rates must also be related (mostly) to business cycles. Furthermore, we find a negative correlation between inflation and unemployment growth and their contemporaneous counterparts, which is the same for the RavenPack data. These results indicate that our sentiment variables are reasonably connected to the underlying macroeconomic variables and thus should contain exploitable information for the purpose of yield curve forecasting.

Concerning the lead-lag structure between macroeconomic variables and their sentiment counterparts, we find that the direction of the correlation between the macroeconomic variables and their lead or lagged sentiment counterpart is the same as the contemporaneous correlation. Furthermore, we observe that contemporaneous inflation is statistically significantly and positively correlated with lead and lagged unemployment sentiment.

Figure 1.1: Sentiment vs. Macroeconomic Variables

This figure shows different macroeconomic variables in blue against their sentiment counterpart in yellow. All time series are standardized. The data cover the period from January 2008 to May 2022.

	$S_{i,t}$	$S_{\pi,t}$	$S_{u,t}$	$S_{i,t+1}$	$S_{\pi,t+1}$	$S_{u,t+1}$	$S_{i,t-1}$	$S_{\pi,t-1}$	$S_{u,t-1}$
Δi_t	**0.22**	0.05	-0.15	**0.17**	0.02	-0.12	**0.26**	0.04	-0.04
π_t	-0.06	**-0.24**	**0.31**	-0.05	**-0.23**	**0.34**	-0.06	**-0.22**	**0.3**
u_t	-0.11	0.08	-0.11	-0.15	0.13	**-0.21**	-0.15	0.05	-0.06

Table 1.1: Correlation Analysis Sentiment vs. Macroeconomic Variables
The table depicts the correlation between the Twitter sentiment variables (contemporaneously, lead one period and lagged one period) and their macroeconomic counterparts. Δi_t, π_t and u_t represent the first difference in the short rate, inflation and unemployment growth, respectively. S_i, S_π, S_u denote the sentiment on each macroeconomic variable. Bold font indicates a p-value lower than 0.05, testing whether the true correlation coefficient is equal to zero. Due to the relatively small number of posts during the initial years of Twitter, the time period ranges from January 2011 to May 2022.

1.3.3 Yield Curve Data

To construct the U.S. Treasury zero-coupon yield curve, we use the factors made available by Gürkaynak et al. (2007). We obtain the time series of yields with maturities of 3-, 12-, 24-, 36-, 48-, 60-, 72-, 84-, 96-, 108-, and 120-months for the time period January 2008 until May 2022. We provide summary statistics for the constant maturity yields in Table 1.2. On average, the yield curve is increasing and concave, consistent with established stylized facts. Furthermore, we see that the standard deviation is on average decreasing with maturity, which is consistent with the idea that future yields contain expected values (thus average) of future short-term yields (Audrino & Offner, 2022). We also document non-normality of the yields in terms of negative excess kurtosis and positive skewness for the yields at all maturities. The finding that the serial correlation on average decreases with maturity differs from the results in Diebold and Li (2006). Finally, in Panel B, we show that the correlations of yields across maturities is high and on average the highest for yields with adjacent maturities, ranging between 0.92 and 0.99.

Panel A								
	3-month	6-month	1-year	2-year	3-year	5-year	7-year	10-year
Mean	0.72	0.72	0.77	0.96	1.19	1.66	2.06	2.49
Median	0.27	0.28	0.35	0.71	1.03	1.65	2.07	2.4
Maximum	2.71	2.66	2.7	2.86	2.92	3.46	3.89	4.84
Minimum	0.03	0.08	0.07	0.11	0.14	0.24	0.37	0.55
Std. Deviation	0.76	0.77	0.78	0.77	0.74	0.74	0.79	0.9
Skewness	1.18	1.14	1.09	0.98	0.72	0.17	0.08	0.21
Kurtosis	2.93	2.81	2.75	2.78	2.65	2.38	2.48	2.6
ACF 1	0.96	0.97	0.96	0.95	0.94	0.94	0.94	0.95

Panel B								
Correlations	3-month	6-month	1-year	2-year	3-year	5-year	7-year	10-year
3-month	1							
6-month	0.99	1						
1-year	0.96	0.99	1					
2-year	0.89	0.93	0.98	1				
3-year	0.81	0.86	0.91	0.98	1			
5-year	0.6	0.64	0.71	0.83	0.92	1		
7-year	0.4	0.44	0.5	0.63	0.76	0.95	1	
10-year	0.22	0.24	0.29	0.41	0.57	0.83	0.96	1

Table 1.2: U.S. Treasury Yield Curve Summary Statistics

This table shows in Panel A summary statistics for the U.S. Treasury yield curve at various maturities. The data are expressed in annualized percentage terms. We use monthly data for the sample period January 2008 - May 2022. Panel B depicts the correlations between pairs of Treasury yields.

In addition, by analyzing first differences of yields, it becomes evident that the volatility is not constant. In Figure 1.2, we depict the path of the first difference of the 3-month and the 10-year yield. For both series, we see spikes in volatility around the Great Financial Crisis (GFC) in 2008. Furthermore, the outbreak of the Coronavirus in spring 2020 led to another visible spike in the volatility of the 3-month yield. Such volatility clusters in the yields have attracted scholarly attention (Audrino, 2006; Audrino & Medeiros, 2010) and additionally justify the implementation of regime-switching models for yield curve modelling.

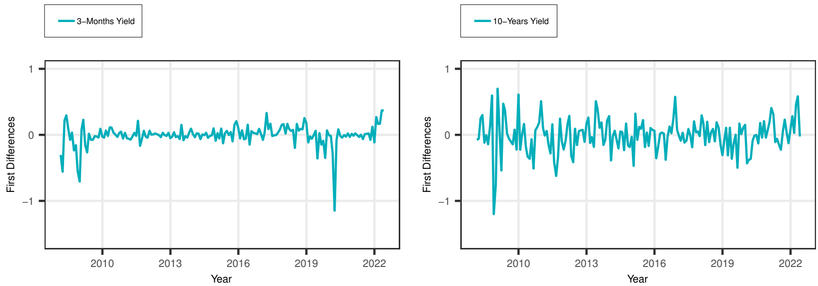

Figure 1.2: Yields First Differences

This figure shows the first differences of 3-month and 10-year U.S. Treasury yields from January 2008 to May 2022.

1.3.4 Monetary policy indicators

Guidolin and Pedio (2019) show that variables capturing the stance of monetary policy add predictive power to the standard DNS factors when forecasting the U.S. Treasury yield curve. We extend our benchmark analysis by also comparing our sentiment-augmented models against the monetary policy indicator-augmented models presented in Guidolin and Pedio (2019). Specifically, we include the same five variables that represent the features of monetary policy adopted by the Fed after the Great Financial Crisis in 2008. Two variables are chosen to represent the size of the Fed's balance sheet. The first is the level of the Fed's total asset, for which we use the data series published on FRED by the Federal Reserve Bank of St. Louis.[2] The second is a divisia money aggregate index. We use the log level of the M3 Index for the United States

[2]The series is found using the label *WALCL*.

obtainable on FRED.[3,4] Next, we include two variables that approximate the composition of the Fed's balance sheet. The first variable describes the percentage of Treasuries on the Fed's total balance sheet. The total amount of securities outstanding in any given month is obtained on FRED. For the analysis, we use the corresponding end of month values and divide them with the size of the total balance sheet.[5] The second variable describing the average maturity of the Fed's portfolio of Treasuries is constructed by taking the sum of the weighted maturities, where the weights correspond to the percentage of Treasuries with a specific maturity of the total amount of Treasuries outstanding. Finally, we use the Fed funds rate to proxy the interest rate policy of the Fed.[6] For a detailed description of how these variables relate to monetary policy, we refer to Guidolin and Pedio (2019).

1.4 Statistical Models

1.4.1 Baseline Model

Our starting point is the standard Nelson and Siegel (1987) model to fit the forward yield curve in a static setting:

$$f(\tau) = \beta_1 + \beta_2 e^{-\tau\lambda} + \beta_3 \lambda \tau e^{-\tau\lambda}, \tag{1.2}$$

[3]The series is labelled *MABMM301USM189S*.

[4]We do not use the MZM Index as in Guidolin and Pedio (2019) since its publication terminated in 2021.

[5]The FRED label for the data series describing the total amount of Treasuries outstanding is *TREAST*.

[6]The data are downloadable from FRED using the label *FEDFUNDS*.

where τ is the maturity, and $\beta_1, \beta_2, \beta_3$ and λ are function parameters. The connection between the instantaneous forward rate and the yield is derived from the no-arbitrage theory and given by:

$$y(\tau) = \frac{1}{\tau} \int_0^\tau f(u)du. \tag{1.3}$$

Equation (1.4) represents the dynamic Nelson-Siegel model (DNS) which obtains from inserting (1.2) into (1.3) and allowing for time dynamics as in Diebold and Li (2006):

$$y_t(\tau) = \beta_{1,t} + \beta_{2,t}\left(\frac{1-e^{-\tau\lambda}}{\tau\lambda}\right) + \beta_{3,t}\left(\frac{1-e^{-\tau\lambda}}{\tau\lambda} - e^{-\tau\lambda}\right). \tag{1.4}$$

$y_t(\tau)$ represents the yield to maturity τ at time t and $\beta_{1,t}$, $\beta_{2,t}$, $\beta_{3,t}$ are the latent factors, which control the dynamics of $y_t(\tau)$ for all maturities. 1, $\left(\frac{1-e^{-\tau\lambda}}{\tau\lambda}\right)$, and $\left(\frac{1-e^{-\tau\lambda}}{\tau\lambda} - e^{-\tau\lambda}\right)$ are the factor loadings. In terms of interpretation of the DNS factors, we follow Diebold and Li (2006) and rename $\beta_{1,t}$, $\beta_{2,t}$, $\beta_{3,t}$ as level factor L_t, slope factor S_t and curvature factor C_t. Since the implied term structure gets steeper the lower the slope factor is, it is actually a (negative) slope factor. Following the established literature — Byrne et al. (2017), Coroneo et al. (2016), Diebold and Li (2006), Guidolin and Pedio (2019) — we fix the decay parameter λ at 0.0609. Consequently, the curvature factor is maximized at exactly 30-months, consistent with the discussion by Yu and Zivot (2011). We use the state-space representation introduced by Diebold et al. (2006) to get $N \cdot T$ measurement equations. Specifically, the vector of measurement equations reads:

$$\begin{pmatrix} y_t(\tau_1) \\ y_t(\tau_2) \\ \vdots \\ y_T(\tau_N) \end{pmatrix} = \begin{pmatrix} 1 & \frac{1-e^{-\tau_1\lambda}}{\tau_1\lambda} & \frac{1-e^{-\tau_1\lambda}}{\tau_1\lambda} - e^{-\tau_1\lambda} \\ 1 & \frac{1-e^{-\tau_2\lambda}}{\tau_2\lambda} & \frac{1-e^{-\tau_2\lambda}}{\tau_2\lambda} - e^{-\tau_2\lambda} \\ \vdots & \vdots & \vdots \\ 1 & \frac{1-e^{-\tau_N\lambda}}{\tau_N\lambda} & \frac{1-e^{-\tau_N\lambda}}{\tau_N\lambda} - e^{-\tau_N\lambda} \end{pmatrix} \begin{pmatrix} L_t \\ S_t \\ C_t \end{pmatrix} + \begin{pmatrix} \epsilon_t(\tau_1) \\ \epsilon_t(\tau_2) \\ \vdots \\ \epsilon_T(\tau_N) \end{pmatrix}, \quad (1.5)$$

where $\epsilon(\tau)$ represents maturity-specific errors, which are IID over time and simultaneously cross-sectionally uncorrelated as in Diebold, Rudebusch, and Aruoba (2006). Fixing the decay parameter λ allows us to estimate (1.5) using OLS. For every point in time t, we fit the riskless yield curve obtaining the three-dimensional time series of DNS factors, $\left\{\hat{L}_t, \hat{S}_t, \hat{C}_t\right\}_t^T$. The latter are used to build a multivariate Gaussian VAR(p) model that constitutes our baseline model. Notationwise, we follow Guidolin and Pedio (2019) and describe the model as:

$$\boldsymbol{f_{t+1}} = \begin{pmatrix} L_{t+1} \\ S_{t+1} \\ C_{t+1} \end{pmatrix} = \begin{pmatrix} \mu_L \\ \mu_S \\ \mu_C \end{pmatrix} + \sum_{j=1}^{p} \begin{pmatrix} \phi_{1,1}^j & \phi_{1,2}^j & \phi_{1,3}^j \\ \phi_{2,1}^j & \phi_{2,2}^j & \phi_{2,3}^j \\ \phi_{3,1}^j & \phi_{3,2}^j & \phi_{3,3}^j \end{pmatrix} \begin{pmatrix} L_{t+1-j} \\ S_{t+1-j} \\ C_{t+1-j} \end{pmatrix} + \begin{pmatrix} \eta_{L,t+1} \\ \eta_{S,t+1} \\ \eta_{C,t+1} \end{pmatrix}, \quad (1.6)$$

which we abbreviate as follows:

$$\boldsymbol{f}_{t+1} = \boldsymbol{\mu} + \sum_{j=1}^{p} \boldsymbol{\Phi}_j \boldsymbol{f}_{t+1-j} + \boldsymbol{\eta}_{t+1}, \quad (1.7)$$

where $\boldsymbol{\mu}$ represents the vector of intercepts, $\boldsymbol{\Phi}_j$ is a 3x3 matrix that determines factor dynamics, and $\boldsymbol{\eta}$ is the innovation vector for which we assume normality such that $\boldsymbol{\eta}_{t+1} \sim N(0, \boldsymbol{\Sigma})$. We do not impose restrictions on $\boldsymbol{\Sigma}$ such that shocks to the DNS factors may exhibit non-zero contemporaneous correlations.

1.4.2 Extension to include sentiment variables and monetary policy indicators

We extend the baseline model in (1.7) to include sentiment variables related to the state of the macroeconomy and variables pertaining to monetary policy as in Guidolin and Pedio (2019). The additional Q variables are collected in a Q x 1 vector times series denoted $\boldsymbol{m}_{t+1}$. The additional variables in $\boldsymbol{m}_{t+1}$ are included in a bi-directional manner, such that they can affect the dynamics of the DNS factors and vice-versa be predicted by the DNS factors. For a discussion, see for example Aguiar-Conraria et al. (2012), who provide evidence for shifting relationships between the shape of the yield curve and macroeconomic determinants. Let $\boldsymbol{x}_{t+1} \equiv \left[\boldsymbol{f}'_{t+1}\boldsymbol{m}'_{t+1}\right]$ denote the entire vector consisting of the DNS factors and the additional Q variables at time $t+1$. The augmented baseline model from equation (1.7) is adjusted to:

$$\boldsymbol{x}_{t+1} = \boldsymbol{c} + \sum_{j=1}^{p} \boldsymbol{A}_j \boldsymbol{x}_{t+1-j} + \boldsymbol{\zeta}_{t+1}, \tag{1.8}$$

where $\boldsymbol{c}$ represents the new vector of intercepts, $\boldsymbol{A}_j$ is a $(3+Q)$ x $(3+Q)$ matrix that governs the dynamics between the DNS factors and the additional variables, and $\boldsymbol{\zeta}$ represents the innovation vector for which we also assume normality such that $\boldsymbol{\zeta}_{t+1} \sim N\left(0, \boldsymbol{\Omega}\right)$. We do not impose restrictions on $\boldsymbol{\Omega}$, such that shocks to the variables can have non-zero contemporaneous correlations.

1.4.3 Markov switching vector autoregressive models

We extend the analysis and allow for regime switches in the analysis since the literature has shown their benefits when forecasting the dynamics of the riskless interest rate curve: see, among others Audrino and Offner (2022), Filipova et al. (2014), Guidolin and Timmermann (2009). For Markov-switching vector autoregressions with K states and p lags, MSVAR(K, p), we assume that there is a latent state-variable $S_{t+1} = 1, ..., K$, which determines the state of the system and thus causes shifts in the parameters of the model. We only allow for one lag $p = 1$ and two states $K = 2$ in order to avoid over-parametrization. Our MSVAR(2,1) model can be written as:

$$\boldsymbol{x}_{t+1} = \boldsymbol{c}_{S+1} + \boldsymbol{A}_{S_{t+1}} \boldsymbol{x}_t + \boldsymbol{\Omega}^{1/2}_{S_{t+1}} \boldsymbol{\nu}_{t+1} \qquad \boldsymbol{\nu}_t \sim IID \quad N(0, \boldsymbol{I}_M), \quad (1.9)$$

where $\boldsymbol{\nu}_{t+1} \equiv \boldsymbol{\Omega}^{-1/2}_{S_{t+1}} \boldsymbol{\zeta}_{t+1}$. Furthermore, note that $\boldsymbol{\Omega}^{1/2}_{S_{t+1}} \left(\boldsymbol{\Omega}^{1/2}_{S_{t+1}}\right)' = Var\left(\boldsymbol{\zeta}_{t+1} \mid S_{t+1}\right)$. Concerning the Markov property, we assume that the latent state variable S_{t+1} is the result of a discrete, time homogeneous, irreducible and ergodic first-order Markov chain such that:

$$Pr\left(S_t = j \mid \{S_j\}^{t-1}_{\tau=1}, \{\boldsymbol{x}_l\}^{t-1}_{l=1}\right) = Pr\left(S_t = j \mid S_{t-1} = i\right) = p_{i,t}, \quad (1.10)$$

where $p_{i,j}$ is the $[i, j]$ element of the K x K transition matrix, $\boldsymbol{P}$ as in (Guidolin & Pedio, 2019). We estimate the MSVAR models using maximum likelihood, applying the blockwise optimization technique introduced by Sims et al. (2008). In order to forecast the yield curve factors in the MSVAR models, we follow Hamilton (1994). The forecasts can be derived as:

$$\hat{\boldsymbol{x}}_{t+1} = \sum_{S=1}^{K} \hat{\xi}_{S_{t+1}} \hat{\boldsymbol{A}}_{S_{t+1}} \boldsymbol{x}_t, \tag{1.11}$$

where $\xi_{S_{t+1}}$ denotes the probability that at $t+1$ state S will prevail. Furthermore, we use $E\left[\boldsymbol{\xi}_{t+1} \mid \boldsymbol{\xi}_t\right] = \boldsymbol{P}' \boldsymbol{\xi}_t$: see Hamilton (1994).

1.5 Trading Strategy Description

In this section, we dig deeper into earlier results and investigate whether and to what extent sentiment-augmented models used to predict the interest term structure can be applied to implement a systematic trading strategy. Specifically, we build bond butterfly strategies and exploit our macro sentiment data. We explain how to build butterfly strategies and illustrate how to measure their economic value.

The choice of bond butterflies as a trading strategy is motivated by their ubiquitous use in practice to profit from views on changes in the shape of the yield curve. We follow Guidolin and Pedio (2019) and Fabozzi et al. (2005), who have shown how to use them within a Nelson-Siegel framework. Specifically, we build semi-hedged, long-short, self-financing strategies in which we are exposed either to the slope or curvature factor, while staying neutral to the other factors. For example, if we predict and bet on a change in the slope of the yield curve, we neutralize our exposure of the portfolio to level and curvature movements. Concretely, we build long and short positions at different maturities to hedge against level and curvature movements of the yield curve. For example, we predict a change in the slope of the yield curve between t and $t+1$, but we have no view concerning the development of the curvature and the level of the curve. We build a duration-weighted hedged and self-financing bond butterfly using three bonds with short (2-years), medium (5-years), and long maturity (10-years). We obtain the portfolio weights, from the following system of equations:

$$\begin{cases} q_s\, MD_s\, L_s + q_m\, MD_m\, L_m + q_l\, MD_l\, L_l = 0, \\ q_s\, MD_s\, C_s + q_m\, MD_m\, C_m + q_l\, MD_l\, C_l = 0, \\ q_s + q_m + q_l = 1, \end{cases} \tag{1.12}$$

where q_s, q_m, and q_l are the weights of the short-, medium-, and long-term bonds, MD_s; MD_m and MD_l are the modified durations of the short-, medium-, and long-term bonds; L_s, L_m and L_l are the sensitivities of the short-, medium-, and long-term yields to the level factor as calculated in Nelson-Siegel's model; and finally, C_s, C_m and C_l, are the exposures of the short-, medium-, and long-term yields to the curvature factor (i.e. the loadings of curvature).

We build butterflies based on 2-, 5-, and 10-year bonds.[7] The direc-

[7]We do not consider other butterfly schemes such as the (2/5/7) or (2/10/30) butterfly in order to avoid multiple testing issues.

tion of the butterfly trades can take two distinct forms: barbell and bullet. In a barbell butterfly, we go long on the wings and short the body. This means that we buy the bonds with 2-years and 10-years maturity and sell bonds with 5-years maturity. In a bullet butterfly we buy the body and sell the wings. Which strategy is executed depends on our yield curve shape forecasts generated by the models described in Section 1.4. We re-estimate the model every month and use a rolling window for the estimation. As in Fabozzi et al. (2005), when we predict an increase (decline) of the slope or curvature factor, we implement the bullet (barbell) butterfly if the sensitivity is positive. The sensitivity of a Nelson-Siegel-weighted butterfly is derived in Martellini et al. (2003, p. 197). The positions are held for a period of one month. As in Fabozzi et al. (2005) we find negative (positive) sensitivities of our butterflies to the slope (curvature) factor.

In order to measure the economic performance of the strategy outlined above, we determine the monthly returns and average them over the entire out-of-sample period. Monthly returns are calculated using the definition of the modified duration of a bond. The percentage return of a given bond between t and $t+1$ is commonly approximated by the negative of the product of the modified duration and the change in the yield between t and $t+1$. The return of a bond $r_{i,t}$ in a given month t is given by:

$$r_{i,t} = -MD_i \cdot \Delta Y_i, \tag{1.13}$$

where MD_i is the modified duration of bond i and ΔY_i represents the change in the yield of bond i in percentage points between period t and $t+1$. The return approximation in equation (1.13) neglects convexity effects on returns. If changes in the yields are small, which they usually are on a monthly level, the approximation is justified. We weight the monthly returns using the weights q obtained in equation (1.12). Hence, the return of the butterfly in a given month can be stated as:

$$R_t = \sum_{j=s,m,l} q_{j,t} r_{j,t} \tag{1.14}$$

where s, m, l represent short-, medium-, and long-term maturity. Furthermore, for the case when we short the butterfly in order to execute a bullet trade, we invest the proceeds for one month at the risk-free interest rate. Our trading period ranges from January 2016 to May 2022. Hence, in total we have 77 monthly butterfly returns, of which we calculate the monthly average performance.

Concerning transaction costs, we follow the literature (Guidolin & Pedio, 2019) and deduct from the monthly returns the costs for funding the margin positions, assumed to be ten percent of the current repo rate. The repo rates are given by the Federal Reserve Bank of New York.

Furthermore, we calculate the Sharpe ratio SR_p of the returns for trading period p, which is given by:

$$SR_p = \frac{\frac{1}{P}\sum_{t=1}^{P} R_t}{\sqrt{\frac{1}{P}\left(\sum_{t=1}^{P} R_t - \frac{1}{P}\sum_{t=1}^{P} R_t\right)^2}}, \tag{1.15}$$

where the numerator does not include the risk-free rate adjustment because the strategy derived in equation (1.12) is zero net-outlay, which means that any funding costs are deducted from the total return; see Guidolin and Pedio (2019).

1.6 Empirical Results

1.6.1 Discussion Term Structure Models

Table 1.3 shows the results of the sentiment data-augmented dynamic Nelson-Siegel model described in equation (1.8). All three factors (level, slope, curvature) are highly persistent, while the degree of persistence varies across factors. In line with stylized facts, the model implies that long-term yields are less persistent than short-term yields. This is because the long-term yields are mainly driven by the least persistent level factor, whereas short-term yields are affected both by the level and the slope factor.

We observe statistically significant autoregressive coefficients at the 1% level for all variables except inflation sentiment, which is significant at the 10% level. Cross-sentiment dynamics is as follows: The sentiment on interest rates influences negatively the sentiment on unemployment and vice versa. Concerning sentiment-factor dynamics, we find that the inflation sentiment positively affects the level factor and that the unemployment sentiment negatively affects the slope factor. Concerning lagged effects of factors on sentiment, we find statistically significant impacts of the lagged level factor on the unemployment sentiment as well as from the lagged curvature factor on current interest and unemployment sentiment.

	Dependent Variables					
	L_t	S_t	C_t	$S_{i,t}$	$S_{\pi,t}$	$S_{u,t}$
L_{t-1}	0.907***	0.075	−0.137	0.170	−0.050	0.245**
	(0.042)	(0.048)	(0.111)	(0.154)	(0.152)	(0.103)
S_{t-1}	−0.057*	0.975***	−0.007	−0.175	−0.028	0.066
	(0.030)	(0.035)	(0.090)	(0.114)	(0.121)	(0.086)
C_{t-1}	0.005	0.042***	0.935***	0.077*	0.051	0.073**
	(0.014)	(0.014)	(0.043)	(0.042)	(0.043)	(0.031)
$S_{i,t-1}$	0.002	−0.004	0.0001	0.216***	−0.027	−0.105*
	(0.032)	(0.028)	(0.082)	(0.083)	(0.097)	(0.058)
$S_{\pi,t-1}$	0.040*	−0.034	−0.067	−0.039	0.140*	−0.092
	(0.023)	(0.025)	(0.052)	(0.096)	(0.084)	(0.075)
$S_{u,t-1}$	0.037	−0.075**	0.058	−0.226**	−0.187	0.588***
	(0.027)	(0.030)	(0.088)	(0.104)	(0.124)	(0.089)
μ	0.192**	−0.136	0.206	−0.691**	0.332	−0.365*
	(0.092)	(0.099)	(0.262)	(0.297)	(0.302)	(0.213)
	Estimated Correlation Matrix					
	L_t	S_t	C_t	$S_{i,t}$	$S_{\pi,t}$	$S_{u,t}$
L_t	-	-	-	-	-	-
S_t	-0.82***	-	-	-	-	-
C_t	-0.30***	-0.01	-	-	-	-
$S_{i,t}$	-0.10	0.14	0.06	-	-	-
$S_{\pi,t}$	-0.09	0.07	0.11	0.09	-	-
$S_{u,t}$	0.12	-0.21**	-0.02	-0.21***	-0.14	-
Observations	172	172	172	172	172	172
Adjusted R^2	0.935	0.951	0.915	0.182	0.055	0.554

Table 1.3: Regression Output of the Dynamic Nelson Siegel Model

The upper table shows the coefficients of the first-order VAR process for the yield curve factors and sentiment from January 2008 to May 2022. The factors are estimated under the Dynamic Nelson-Siegel model. HC3 standard errors are shown in parentheses. The middle table reports the residual correlation matrix. Inference is made under the assumption of white-noise residuals. ***, **, and * denote significance at the 1 5 and 10 level, respectively. The adjusted R squared describes the model fit for the respective dependent variable.

In Table 1.4 and 1.5 the coefficients of the MSVAR model from equation (1.9) are depicted. The first (second) table depicts the coefficients of regime 1 (2). We observe statistically significant autoregressive coefficients for the DNS at the 1% level. Regime 1 coefficients for the DNS factors are very similar to those found in the simple VAR model in Table 1.3. However, in regime 2, we observe a much lower level factor, which determines yield level dynamics. Thus, regime 2 exhibits lower serial persistency with respect to the yield level. Consistent with this finding, Figure **??** reveals that regime 2 is instigated, *inter alia*, when there are abrupt changes in the yield levels as during the Great Financial Crisis or as during the onset of the COVID-19 pandemic in March 2020.[8] Since we are estimating a MSVAR model not only for the DNS factors, the shape factors of the yield curve, but also for sentiment variables, changes to the sentiment variables can also induce regime shifts. The tendency towards regime 2 during the year 2021 can be explained by lower levels of inflation sentiment. The interest rate sentiment and unemployment sentiment do not show any noticeable regime-shifting dynamics during that period, as shown in Figure 1.1. However, inflation sentiment reverts to the mean in 2022, while regime 2 still dominates during 2022, which defies interpretation.

Furthermore, we find that sentiment variables have a larger impact on the DNS factors and thus the shape of the yield curve in regime 2, as depicted in Table 1.5. For example, in regime 2, a one standard deviation increase in lagged interest rate sentiment statistically significantly decreases the level and increases the slope and curvature factor by -0.177, 0.094, 0.633, respectively. However, no impact of lagged interest rate sentiment on the yield curve factors was found in regime 1. Similarly, a one standard deviation increase in lagged inflation sentiment increases the level and decreases the slope and curvature factor by 0.270, -0.242, -0.442, respectively. In regime 1, lagged inflation sentiment has no statistically significant impact on the DNS factors. Finally, lagged unem-

[8] We use the Fed Funds Rate for illustrative purposes as it impacts Treasury yields at all maturities.

ployment sentiment statistically significantly influences level and slope dynamics in regime 1. While the coefficients are statistically significant, they are comparably small in economic terms. However, in regime two, lagged unemployment sentiment impacts the curvature factor statistically and economically significantly. A one standard deviation increase in lagged unemployment sentiment increases the curvature factor by 0.233. Consequently, we argue that sentiment variables matter, especially in regime 2, which characterizes times of fast changing yield or inflation dynamics.

Moreover, all statistically significant coefficients in regime 1 describing the lagged factor-to-sentiment and cross-sentiment dynamics have the same sign as the corresponding coefficients from the simple VAR model depicted Table 1.3. Interestingly, in regime 2, we observe numerous deviations in terms of the direction of the effect compared to regime 1 for those coefficients. For example, while in regime 1 lagged unemployment sentiment negatively and statistically significantly correlates with current interest and inflation sentiment, the impact is positive and statistically significant in regime 2. In other words, not only sentiment-to-factor dynamics but also cross-sentiment dynamics show regime dependence.

Finally, looking at the transition matrix depicted in Table 1.4, we see that regime 1 is more persistent compared to regime 2. The probability of staying in regime 1 is 98.1%, while the probability of staying in regime 2 is 89.3%. Consistently, we observe 115 regime 1 and only 57 regime 2 observations with an average duration of 38.3 months for regime 1 and 19 months for regime 2.

Panel A	*Regime 1: Coefficients*					
	L_t	S_t	C_t	$S_{i,t}$	$S_{\pi,t}$	$S_{u,t}$
L_{t-1}	0.905***	0.085*	0.052	0.15	−0.172	0.279***
	(0.043)	(0.048)	(0.119)	(0.154)	(0.154)	(0.102)
S_{t-1}	−0.025	0.972***	0.058	−0.226**	−0.17	0.075
	(0.031)	(0.035)	(0.094)	(0.112)	(0.122)	(0.085)
C_{t-1}	−0.014	0.047***	0.945***	0.123***	0.115**	0.086***
	(0.016)	(0.014)	(0.044)	(0.044)	(0.05)	(0.031)
$S_{i,t-1}$	0.032	−0.026	−0.111	0.247***	−0.07	−0.126**
	(0.035)	(0.028)	(0.087)	(0.084)	(0.102)	(0.057)
$S_{\pi,t-1}$	0.025	−0.008	−0.001	0.106	0.106	−0.149**
	(0.023)	(0.026)	(0.055)	(0.091)	(0.088)	(0.073)
$S_{u,t-1}$	0.049*	−0.066**	0.007	−0.284**	−0.284**	0.585***
	(0.029)	(0.03)	(0.084)	(0.107)	(0.128)	(0.089)
μ	0.189*	−0.132	−0.26	0.702**	0.702**	−0.423*
	(0.095)	(0.101)	(0.277)	(0.302)	(0.318)	(0.213)

	Regime 1: Error Correlation Matrix					
	L_t	S_t	C_t	$S_{i,t}$	$S_{\pi,t}$	$S_{u,t}$
L_t	-	-	-	-	-	-
S_t	-0.81***	-	-	-	-	-
C_t	-0.31***	-0.01	-	-	-	-
$S_{i,t}$	-0.09	0.14	0.04	-	-	-
$S_{\pi,t}$	-0.08	0.07	0.08	0.12	-	-
$S_{u,t}$	0.11	-0.21***	-0.02	-0.21***	-0.14	-

Estimated Transition Matrix		
	Regime 1	*Regime 2*
Regime 1	0.981	0.019
Regime 2	0.107	0.893

Table 1.4: Regression Output of MSVAR(1,2) Sentiment Model: Panel A

The upper table shows the coefficients of the MSVAR(1,2) model for the yield curve factors and sentiment from January 2008 to May 2022. The factors are estimated under the Dynamic Nelson-Siegel model. HC3 standard errors are shown in parentheses. The middle table reports the residual correlation matrix. Inference is made under the assumption of white-noise residuals. ***, **, and * denote significance at the 1%, 5% and 10% level, respectively. The lowest table reports the regime transition matrix.

Panel B	*Regime 2: Coefficients*					
	L_t	S_t	C_t	$S_{i,t}$	$S_{\pi,t}$	$S_{u,t}$
L_{t-1}	0.643***	0.083	−0.219	0.073	0.846***	−0.433*
	(0.095)	(0.064)	(0.192)	(0.204)	(0.272)	(0.252)
S_{t-1}	−0.489***	0.962***	0.478***	−0.311*	1.047***	−0.653***
	(0.09)	(0.049)	(0.165)	(0.174)	(0.241)	(0.194)
C_{t-1}	0.047*	0.058***	0.909***	−0.063	0.13*	−0.174**
	(0.025)	(0.019)	(0.053)	(0.06)	(0.071)	(0.083)
$S_{i,t-1}$	−0.177***	0.094**	0.633***	−0.136	0.396**	−0.271*
	(0.044)	(0.04)	(0.128)	(0.112)	(0.167)	(0.143)
$S_{\pi,t-1}$	0.27***	−0.242***	−0.442***	−0.156	−0.156	0.513***
	(0.056)	(0.043)	(0.092)	(0.168)	(0.162)	(0.193)
$S_{u,t-1}$	0.074	−0.037	0.233*	0.594***	0.594***	−0.106
	(0.046)	(0.041)	(0.139)	(0.127)	(0.205)	(0.173)
μ	0.032	−0.3**	1.778***	−0.33	−0.33	0.079
	(0.145)	(0.127)	(0.395)	(0.36)	(0.482)	(0.562)
	Regime 2: Error Correlation Matrix					
	L_t	S_t	C_t	$S_{i,t}$	$S_{\pi,t}$	$S_{u,t}$
L_t	-	-	-	-	-	-
S_t	-0.76***	-	-	-	-	-
C_t	-0.35***	0.13	-	-	-	-
$S_{i,t}$	-0.36***	0.29***	0.13	-	-	-
$S_{\pi,t}$	-0.35***	0.12	0.13	0.20***	-	-
$S_{u,t}$	0.18**	-0.25***	-0.06	-0.21***	-0.11	-
	Regime Duration Information					
		Regime 1		*Regime 2*		
# Obs.		115		57		
Dur. Aver.		38.3		9		
N	172	172	172	172	172	172
Ad. R^2	0.946	0.956	0.928	0.261	0.168	0.653

Table 1.5: Regression Output of MSVAR Sentiment Model: Panel B

The upper table shows the coefficients of the MSVAR(2,1) model for the yield curve factors and sentiment from January 2008 to May 2022. The factors are estimated under the Dynamic Nelson-Siegel model. HC3 standard errors are shown in parentheses. The middle table reports the residual correlation matrix. Inference is made under the assumption of white-noise residuals. ***, **, and * denote significance at the 1%, 5% and 10% level, respectively. The lowest table depicts regime duration information. Regime prevalence is assumed if its filtered probability is higher than 50%. The adjusted R squared describes the model fit for the respective dependent variable over the entire sample, using regime 1 and regime 2.

1.6.2 Discussion trading results

The trading results from betting on changes of the slope factor (curvature factor) are depicted in Table 1.6 (1.7). The results are split into different lengths of trading periods. Panel *A* describes the results for the entire trading period from January 2016 to May 2022, while Panel *B* describes the time period before the outbreak of the Corona pandemic ending in February 2020, and Panel *C* includes the trading results for the Corona time period starting in March 2020. There are 50 trading months before the Corona pandemic and 27 trading months during Corona. In each panel, we show the top six and bottom six models out of 30 models with respect to cumulative return, separated by the dashed line in each panel.[9] The complete list of models is depicted in Appendix 1.C. In order to benchmark the cumulative returns of our trading strategy, we build an equally weighted, long only benchmark portfolio of 2-year-, 5-year-, and 10-year U.S. Treasury notes. The performance is given in Table 1.8.

The results from betting on changes in the slope are depicted in Table 1.6. A total of 9 models beat the benchmark in terms of cumulative return before Corona, while all butterfly strategies beat the benchmark during Corona. The gains from the butterfly strategy during the Corona period lift the results for the entire trading period, such that all 30 models beat the benchmark portfolio concerning the entire trading period. The top six models all have Sharpe ratios higher than 1 in all trading

[9] The numerical results of the other models are available from the authors upon request.

periods shown in panels A, B, C. Furthermore, the p-values of the T-tests of the bottom six models indicate that at the 5% level, we can only statistically significantly differentiate two (six, two) models from the best performing model in terms of their mean return during the respective trading period depicted in Panel A (B, C). Similarly, the bootstrapped p-values of the Sharpe ratios of the bottom six models show that we can statistically significantly differentiate five (four, one) from the best performing model in terms of the Sharpe ratio, as illustrated in Panel A (B, C). With respect to minimum cumulative return, there are some models that never lost money over the entire trading period, which does imply that those models are the best in terms of cumulative return. In terms of maximum drawdown, the models do not show any notable differences for all three trading periods. The answer to the question of whether macro sentiment data-augmented models outperform other models in terms of realized cumulative return is no. While there are five (three) sentiment-augmented models within the top six in Panel A and B (C), even the best performing model in each category does not statistically significantly outperform the raw models, in which we only use the DNS factors to predict changes in the slope of the yield curve. In Table 1.12, we can see that the model *MSVAR DNS Only (VAR DNS Only)* ranks on place 8 (18), 6 (22), 25 (4) for the entire, before Corona and during Corona periods, respectively, with T-test p-values never below 0.06. Furthermore, we observe that the models are unstable in that well performing models before Corona had much worse results during Corona and vice versa. The benchmark models augmented with monetary policy indicators do not perform significantly better or worse compared to the sentiment-augmented models. To summarize, for slope bets, we do not find convincing evidence that sentiment-augmented models yield better results compared to non-sentiment-augmented models. However, we find many models that yield butterfly trading results that perform consistently better than the benchmark.

The results from the bets on the curvature factor are shown in Table 1.7. For the entire trading period, the top 9 models beat the benchmark in terms of cumulative return. Before Corona, only two models are better than the benchmark and during Corona 15 models beat the benchmark. Overall, the returns from the curvature bets are much more volatile, which results in significantly lower Sharpe ratios in all trading periods compared to the slope bets in Table 1.6. For all models, the Sharpe ratios are lower than 1, except for the model *MSVAR Percentage* during Corona. Again, from a statistical point of view, sentiment-augmented models do not perform significantly better than the raw DNS factor models. In Table 1.12 we show that only during Corona does the best performing model, *MSVAR Percentage*, outperform the raw models at the five percent level. However, this model performs very badly for the 50 trading months before Corona, ranking at place 27. No sentiment-augmented model outperforms the raw models from a statistical point of view. However, in terms of model rank consistency, we find that the model *MSVAR Inflation Sentiment (Ravenpack)* performs consistently very well. For the trading period before Corona, the model ranks at place 5, and during Corona at place 6. For the entire trading period, the model is in third place. While we cannot find statistical evidence that the model performs better than the raw models, economically it delivers significantly better results, beating the benchmark portfolio especially during Corona but not before Corona. Interestingly, for every trading period, there are at least two inflation sentiment-augmented models within the top six models. This delivers some evidence that inflation sentiment can be helpful in predicting changes in the curvature of the yield curve, which can be translated into real economic returns. In particular, the cumulative returns are high for the time period during Corona. In fact, inflation has been sharply rising since the onset of the Corona pandemic and the yield curve has experienced decreasing levels of curvature. This is consistent with the highly negative and statistically and economically significant coefficient describing the relationship between lagged inflation sentiment and the curvature factor in regime 2, which dominates during the Corona trading period.

	Cumulative Return (in %)	Average Monthly Return (in %)	T-Test p-value	Sharpe Ratio	SR p-value	Max. Drawdown (in %)	Min. Cum. Return (in %)	Max. Cum. Return (in %)	R.B.C.	R.D.C.
Panel A: Entire Out-of-Sample Period										
VAR Michigan Cons. Sentiment	10.84	0.13	1	1.74	1	1.04	-1.04	10.99	4	10
VAR All Sent. (Ravenpack)	10.7	0.13	0.97	1.72	0.95	0.9	-0.9	10.85	3	13
MSVAR Michigan Cons. Sentiment	10.03	0.12	0.83	1.61	0.83	0.74	0	10.51	1	19
MSVAR Unemp. Sent. (Twitter)	9.84	0.12	0.79	1.58	0.74	0.57	0	10.34	2	21
VAR Unemp. Sent. (Ravenpack)	9.28	0.12	0.67	1.48	0.25	0.9	-0.9	9.43	7	15
VAR Fundsrate	8.34	0.1	0.5	1.33	0.31	1.17	0	8.49	9	12
MSVAR All Sent. (Twitter)	4.38	0.06	0.08	0.69	0.1	2.25	-0.73	6.57	11	30
VAR Percentage	4.34	0.06	0.07	0.71	0.02	1.04	-1.04	4.48	30	2
MSVAR Maturity	4.24	0.05	0.07	0.67	0.04	1.88	-1.88	5.04	21	23
VAR Inf. Sent. (Ravenpack)	4.07	0.05	0.06	0.66	0	1.21	-1.04	4.22	29	14
MSVAR Log Assets	3.7	0.05	0.05	0.59	0.03	2.51	-2.51	3.99	28	18
VAR Inf. Sent. (Twitter)	3.22	0.04	0.04	0.52	0	1.15	-1.04	3.36	27	24
Panel B: Before Corona										
	Cumulative Return (in %)	Average Monthly Return (in %)	T-Test p-value	Sharpe Ratio	SR p-value	Max. Drawdown (in %)	Min. Cum. Return (in %)	Max. Cum. Return (in %)	R.A.P.	R.D.C.
MSVAR Michigan Cons. Sentiment	7.61	0.15	1	1.92	1	0.65	0	7.89	3	19
MSVAR Unemp. Sent. (Twitter)	7.59	0.15	0.99	1.92	1	0.57	0	7.87	4	21
VAR All Sent. (Ravenpack)	7.28	0.14	0.91	1.8	0.87	0.9	-0.9	7.55	2	13
VAR Michigan Cons. Sentiment	7.08	0.14	0.85	1.72	0.8	1.04	-1.04	7.35	1	10
MSVAR Inf. Sent. (Twitter)	6.98	0.14	0.83	1.72	0.79	0.73	-0.73	7.25	9	27
MSVAR DNS Only	6.39	0.12	0.67	1.58	0.54	0.57	0	6.67	8	25
VAR Log Divisia	2.04	0.04	0.05	0.5	0.05	1.07	-1.04	2.3	16	1
VAR Int. Sent. (Ravenpack)	1.97	0.04	0.05	0.5	0.09	0.92	0	2.23	21	11
VAR Inf. Sent. (Twitter)	1.29	0.03	0.03	0.33	0.02	1.15	-1.04	1.65	30	24
MSVAR Log Assets	1.08	0.02	0.02	0.27	0.06	2.51	-2.51	1.34	29	18
VAR Inf. Sent. (Ravenpack)	0.86	0.02	0.02	0.22	0.02	1.21	-1.04	1.12	28	14
VAR Percentage	0.49	0.01	0.01	0.13	0.02	1.04	-1.04	0.93	26	2
Panel C: During Corona										
	Cumulative Return (in %)	Average Monthly Return (in %)	T-Test p-value	Sharpe Ratio	SR p-value	Max. Drawdown (in %)	Min. Cum. Return (in %)	Max. Cum. Return (in %)	R.A.P.	R.B.C.
VAR Log Divisia	3.87	0.14	1	1.96	1	0.35	0	4.01	16	25
VAR Percentage	3.83	0.14	0.98	1.94	0.98	0.31	0	3.97	26	30
MSVAR Unemp. Sent. (Ravenpack)	3.53	0.13	0.86	1.73	0.5	0.38	0	3.53	11	17
VAR DNS Only	3.51	0.13	0.85	1.74	0.31	0.38	0	3.65	18	22
VAR All Sent. (Twitter)	3.51	0.13	0.85	1.74	0.31	0.38	0	3.65	7	13
VAR Int. Sent. (Twitter)	3.51	0.13	0.85	1.74	0.31	0.38	0	3.65	17	20
MSVAR DNS Only	1.78	0.07	0.29	0.83	0.1	0.78	0	2.56	8	6
MSVAR Inf. Sent. (Ravenpack)	1.07	0.04	0.16	0.5	0.1	0.93	-0.05	1.94	24	18
MSVAR Inf. Sent. (Twitter)	0.98	0.04	0.15	0.46	0.09	0.86	0	1.84	9	5
MSVAR Int. Sent. (Ravenpack)	0.94	0.04	0.15	0.44	0.19	1.09	-0.67	1.72	22	14
MSVAR Percentage	-0.34	-0.01	0.04	-0.15	0.1	1.05	-0.62	0.43	23	8
MSVAR All Sent. (Twitter)	-0.36	-0.01	0.04	-0.16	0.03	2.14	-0.41	1.73	25	11

Table 1.6: Trading Results Slope Bets With Transaction Costs

This table shows the top six and bottom six models in terms of realized cumulative return over the respective trading period. The top and bottom models are separated with the dashed line in each trading period. Panel A depicts the results covering the period from January 2016- May 2022, Panel B shows the results for the trading period between January 2016 and February 2020, while Panel C includes the results for the period between March 2020 and May 2022. The column labelled T-Test p-value describes the p-value of the t-test of the mean return of the best performing model and the respective row model. The column SR p-value describes the p-value testing equality of the Sharpe ratio of the best performing model and the respective row model (Ledoit & Wolf, 2008). The column maximum drawdown describes the maximum decline in percentage points from a historical peak in terms of cumulative return. The minimum (maximum) cumulative return describes the lowest (highest) cumulative return achieved over the entire trading period. The last two columns describe the rank of the respective model in different trading periods. R.B.C. stands for rank before Corona, and R.D.C. refers to rank during corona, while R.A.P. is the abbreviation for rank in all periods.

	Cumulative Return (in %)	Average Monthly Return (in %)	T-Test p-value	Sharpe Ratio	SR p-value	Max. Drawdown (in %)	Min. Cum. Return (in %)	Max. Cum. Return (in %)	R.B.C.	R.D.C.
Panel A: Entire Out-of-Sample Period										
MSVAR Percentage	6.93	0.09	1	0.29	1	7.43	-7.43	8.67	27	1
VAR Percentage	6.75	0.09	0.99	0.29	0.99	9.1	-5.33	8.49	12	2
MSVAR Inf. Sent. (Ravenpack)	4.97	0.07	0.89	0.22	0.87	13.3	-3.94	9.36	5	6
MSVAR Inf. Sent. (Twitter)	2.65	0.04	0.76	0.13	0.72	15.06	-9.91	5.15	20	3
VAR Unemp. Sent. (Ravenpack)	0.1	0.01	0.63	0.02	0.6	8.96	-3.9	5.06	16	5
VAR Michigan Cons. Sentiment	-0.07	0	0.62	0.02	0.6	11.8	-0.25	11.73	10	8
VAR Int. Sent. (Ravenpack)	-15.2	-0.21	0.09	-0.67	0.15	15.2	-15.2	0	25	17
MSVAR All Sent. (Ravenpack)	-15.64	-0.21	0.08	-0.69	0.18	16.86	-15.64	1.22	13	28
MSVAR Log Divisia	-15.82	-0.22	0.08	-0.71	0.03	18.35	-18.35	0	30	9
MSVAR Int. Sent. (Ravenpack)	-17.49	-0.24	0.05	-0.79	0.1	18.79	-18.79	0	26	23
VAR All Sent. (Ravenpack)	-19.45	-0.27	0.03	-0.91	0.11	21.95	-19.45	2.5	22	29
VAR Maturity	-22.72	-0.33	0.02	-1.1	0.07	25.61	-22.72	2.89	21	30

	Cumulative Return (in %)	Average Monthly Return (in %)	T-Test p-value	Sharpe Ratio	SR p-value	Max. Drawdown (in %)	Min. Cum. Return (in %)	Max. Cum. Return (in %)	R.A.P.	R.D.C.
Panel B: Before Corona										
VAR Inf. Sent. (Ravenpack)	9.29	0.18	1	0.76	1	2.91	0	11.43	9	22
MSVAR Unemp. Sent. (Twitter)	5.53	0.11	0.68	0.46	0.39	4.28	-0.13	8.15	12	19
VAR Log Divisia	5.15	0.1	0.64	0.43	0.36	4.89	-0.64	9.31	10	16
MSVAR All Sent. (Twitter)	3.94	0.08	0.55	0.33	0.5	4.46	0	6.5	8	11
MSVAR Inf. Sent. (Ravenpack)	3.91	0.08	0.55	0.33	0.44	5.52	0	9.36	3	6
VAR Int. Sent. (Twitter)	3.38	0.07	0.51	0.29	0.24	4.81	-0.47	7.46	13	18
VAR Int. Sent. (Ravenpack)	-5.25	-0.1	0.09	-0.44	0.13	6.56	-6.56	0	25	17
MSVAR Int. Sent. (Ravenpack)	-5.74	-0.11	0.08	-0.48	0.16	5.8	-5.8	0	28	23
MSVAR Percentage	-5.93	-0.12	0.07	-0.5	0.11	7.43	-7.43	0	1	1
MSVAR Unemp. Sent. (Ravenpack)	-8.3	-0.17	0.03	-0.73	0.08	11.32	-11.32	0	15	7
MSVAR Int. Sent. (Twitter)	-8.43	-0.17	0.03	-0.74	0.05	9.49	-9.49	0	20	12
MSVAR Log Divisia	-12.34	-0.26	0.01	-1.16	0.03	12.95	-12.95	0	27	9

	Cumulative Return (in %)	Average Monthly Return (in %)	T-Test p-value	Sharpe Ratio	SR p-value	Max. Drawdown (in %)	Min. Cum. Return (in %)	Max. Cum. Return (in %)	R.A.P.	R.B.C.
Panel C: During Corona										
MSVAR Percentage	13.67	0.49	1	1.2	1	2.17	-1.55	15.53	1	27
VAR Percentage	7.42	0.28	0.59	0.66	0.42	4.78	-4.74	9.17	2	12
MSVAR Inf. Sent. (Twitter)	5.79	0.22	0.5	0.52	0.35	7.16	-7.16	7.51	4	20
MSVAR Michigan Cons. Sentiment	3.91	0.15	0.4	0.36	0.23	5.01	-1.58	5.73	7	24
VAR Unemp. Sent. (Ravenpack)	1.72	0.07	0.3	0.17	0.27	5.6	-2.34	6.69	5	16
MSVAR Inf. Sent. (Ravenpack)	1.01	0.05	0.27	0.11	0.13	7.59	-7.56	2.66	3	5
VAR Unemp. Sent. (Twitter)	-13.05	-0.51	0.01	-1.27	0.09	15.5	-13.05	2.46	23	18
VAR All Sent. (Twitter)	-13.29	-0.52	0.01	-1.31	0.09	14.63	-13.29	1.34	24	17
VAR Inf. Sent. (Twitter)	-14.52	-0.57	0.01	-1.46	0.05	15.99	-15.95	0.03	21	8
MSVAR All Sent. (Ravenpack)	-15.05	-0.59	0.01	-1.53	0.07	15.82	-15.05	0.76	26	13
VAR All Sent. (Ravenpack)	-15.9	-0.63	0	-1.65	0.07	15.9	-15.9	0	29	22
VAR Maturity	-20.11	-0.82	0	-2.35	0.03	20.14	-20.11	0.03	30	21

Table 1.7: Trading Results Curvature Bets With Transaction Costs

This table shows the top six and bottom six models in terms of realized cumulative return over the respective trading period. The top and bottom models are separated with the dashed line in each trading period. Panel A depicts the results covering the period from January 2016- May 2022, Panel B shows the results for the trading period between January 2016 and February 2020, while Panel C includes the results for the period between March 2020 and May 2022. The column labelled T-Test p-value describes the p-value of the t-test of the mean return of the best performing model and the respective row model. The column SR p-value describes the p-value testing equality of the Sharpe ratio of the best performing model and the respective row model (Ledoit & Wolf, 2008). The column maximum drawdown describes the maximum decline in percentage points from a historical peak in terms of cumulative return. The minimum (maximum) cumulative return describes the lowest (highest) cumulative return achieved over the entire trading period. The last two columns describe the rank of the respective model in different trading periods. R.B.C. stands for rank before Corona, and R.D.C. refers to rank during corona, while R.A.P. is the abbreviation for rank in all periods.

	Benchmark Portfolio					
	Cumulative Return (in %)	Average Monthly Return (in %)	Sharpe Ratio	Max Drawdown (in %)	Min. Cum. Return (in %)	Max. Cum. Return (in %)
Entire Out-of-Sample Period	-4.88	-0.06	-0.451	14.44	-5.63	9.21
Before Corona	5.37	0.11	-0.05	10.04	-5.63	5.37
During Corona	-9.73	-0.37	-1.16	13.7	-10.06	3.64

Table 1.8: Benchmark Statistics

The table shows the performance of our benchmark portfolio. The full sample ranges from Janaury 2016 to May 2022. The period before Corona dates from January 2016 until February 2020, while the Corona period starts in March 2020 and ends in May 2022. The benchmark portfolio consists of equally weighted U.S. 2-year-, 5-year-, and 10-year Treasury notes.

1.7 Conclusion

While there exists substantial evidence that sentiment-augmented models predict the yield curve better than non-sentiment-augmented models, we fail to construct trading strategies based upon sentiment data that perform consistently and statistically significantly better than their non-sentiment-augmented counterparts. However, we find mild evidence that inflation sentiment-augmented models perform consistently economically better when betting on changes in the curvature of the yield curve compared to other sentiment-augmented and non-sentiment-augmented models. In particular, the Markov-switching VAR model augmented with inflation sentiment data from RavenPack performs consistently well in all trading periods before as well as during the Corona pandemic. Overall, however, the economic returns from betting on changes in the slope of the yield curve, as opposed to the curvature, are better. Considering the entire trading period, slope bets based on all models applied achieve cumulative returns that are higher than the long-only Treasury note benchmark portfolio. This is due to the better economic returns of the butterfly trading strategy for all underlying models during the Corona pandemic compared to the benchmark. Considering the economic returns of betting on changes in the curvature of the yield curve, the returns are much more volatile, resulting in lower Sharpe ratios and higher draw-downs. Over the entire trading period, only around a third of the models beat the long-only benchmark portfolio. Our models are all Nelson-Siegel type models, which differ from arbitrage-free term structure models.

Testing whether trading strategies based upon sentiment-augmented, arbitrage-free term structure models yield different results presents an interesting future research opportunity.

References

Aguiar-Conraria, L., Martins, M. M., & Soares, M. J. (2012). The yield curve and the macro-economy across time and frequencies. *Journal of Economic Dynamics and Control*, *36*(12), 1950–1970.

Algaba, A., Ardia, D., Bluteau, K., Borms, S., & Boudt, K. (2020). Econometrics meets sentiment: An overview of methodology and applications. *Journal of Economic Surveys*, *34*(3), 512–547.

Audrino, F. (2006). Tree-structured multiple regimes in interest rates. *Journal of Business & Economic Statistics*, *24*(3), 338–353.

Audrino, F., & Medeiros, M. C. (2010). Modeling and forecasting short-term interest rates: The benefits of smooth regimes, macroeconomic variables, and bagging. *Journal of Applied Econometrics*, *26*(6), 999–1022.

Audrino, F., & Offner, E. (2022). The impact of macroeconomic news sentiment on interest rates. *SSRN*, 72. (Available at SSRN: https://ssrn.com/abstract=4170954)

Audrino, F., Sigrist, F., & Ballinari, D. (2020). The impact of sentiment and attention measures on stock market volatility. *International Journal of Forecasting*, *36*(2), 334–357.

Ballinari, D., & Behrendt, S. (2021). How to gauge investor behavior? A comparison of online investor sentiment measures. *Digital Finance*, *3*(2), 169–204.

Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, *2*(1), 1–8.

Byrne, J. P., Cao, S., & Korobilis, D. (2017). Forecasting the term structure of government bond yields in unstable environments. *Journal of Empirical Finance*, *44*, 209–225.

Coroneo, L., Giannone, D., & Modugno, M. (2016). Unspanned macroeconomic factors in the yield curve. *Journal of Business & Economic Statistics*, *34*(3), 472–485.

Daniel, K., Hirshleifer, D., & Teoh, S. H. (2002). Investor psychology in capital markets: evidence and policy implications. *Journal of Monetary Economics*, *49*(1), 139–209.

Diebold, F. X., & Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, *130*(2), 337–364.

Diebold, F. X., Rudebusch, G. D., & Aruoba, S. B. (2006). The macroeconomy and the yield curve: a dynamic latent factor approach. *Journal of Econometrics*, *131*(1-2), 309–338.

Duffee, G. R. (2011). Information in (and not in) the term structure. *Review of Financial Studies*, *24*(9), 2895–2934.

Edison, H. J. (1997). The reaction of exchange rates and interest rates to news releases. *International Journal of Finance & Economics*, *2*(2), 87–100.

Erlwein-Sayer, C. (2017). *Forecasting sovereign bond spreads with macroeconomic news sentiment* (Tech. Rep.). OptiRisk Systems.

Erlwein-Sayer, C. (2018). Macroeconomic news sentiment: Enhanced risk assessment for sovereign bonds. *Risks*, *6*(4), 141.

Fabozzi, F. J., Martellini, L., & Priaulet, P. (2005). Predictability in the shape of the term structure of interest rates. *The Journal of Fixed Income*, *15*(1), 40–53.

Filipova, K., Audrino, F., & Giorgi, E. D. (2014). Monetary policy regimes: Implications for the yield curve and bond pricing. *Journal of Financial Economics*, *113*(3), 427–454.

Goldberg, L., & Grisse, C. (2013). *Time variation in asset price responses to macro announcements* (Tech. Rep.).

Goldberg, L., & Leonard, D. (2005). What moves sovereign bond markets? The effects of economic news on U.S. and German yields. *Federal Reserve Bank of New York*, 1-7. (Available at SSRN: https://ssrn.com/abstract=683269)

Gotthelf, N., & Uhl, M. W. (2018). News sentiment: A new yield curve factor. *Journal of Behavioral Finance*, *20*(1), 31–41.

Guidolin, M., & Pedio, M. (2019). Forecasting and trading monetary policy effects on the riskless yield curve with regime switching Nelson–Siegel models. *Journal of Economic Dynamics and Control*, *107*, 103723.

Guidolin, M., & Timmermann, A. (2009). Forecasts of US short-term interest rates: A flexible forecast combination approach. *Journal of Econometrics*, *150*(2), 297–311.

Gürkaynak, R. S., Sack, B., & Wright, J. H. (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, *54*(8), 2291–2304.

Hamilton. (1994). *Time series analysis*. Princeton: Princeton University Press.

Hevia, C., Gonzalez-Rozada, M., Sola, M., & Spagnolo, F. (2014). Estimating and forecasting the yield curve using a Markov switching dynamic Nelson and Siegel model. *Journal of Applied Econometrics*, *30*(6), 987–1009.

Ledoit, O., & Wolf, M. (2008). Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance*, *15*(5), 850–859.

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, *66*(1), 35–65.

Lucca, D., & Trebbi, F. (2009). *Measuring central bank communication: An automated approach with application to FOMC statements* (Tech. Rep.).

Martellini, L., Priaulet, P., & Priaulet, S. (2003). *Fixed-income securities*. Wiley.

Nelson, C., & Siegel, A. (1987). Parsimonious modeling of yield curves. *The Journal of Business*, *60*(4), 473–489.

Nofer, M., & Hinz, O. (2015). Using Twitter to predict the stock market. *Business & Information Systems Engineering*, *57*(4), 229–242.

Oliveira, N., Cortez, P., & Areal, N. (2013). On the predictability of stock market behavior using StockTwits sentiment and posting volume. In *Progress in Artificial Intelligence* (pp. 355–365). Springer Berlin Heidelberg.

Renault, T. (2017). Intraday online investor sentiment and return patterns in the U.S. stock market. *Journal of Banking & Finance*, *84*, 25–40.

Renault, T. (2019). Sentiment analysis and machine learning in finance: a comparison of methods and models on one million messages. *Digital Finance*, *2*(1-2), 1–13.

Sims, C. A., Waggoner, D. F., & Zha, T. (2008). Methods for inference in large multiple-equation Markov-switching models. *Journal of Econometrics*, *146*(2), 255–274.

Sul, H. K., Dennis, A. R., & Yuan, L. I. (2016). Trading on twitter: Using social media sentiment to predict stock returns. *Decision Sciences*, *48*(3), 454–488.

Sun, L., Najand, M., & Shen, J. (2016). Stock return predictability and investor sentiment: A high-frequency perspective. *Journal of Banking & Finance*, *73*, 147–164.

Tan, S. D., & Tas, O. (2020). Social media sentiment in international stock returns and trading activity. *Journal of Behavioral Finance*, *22*(2), 221–234.

Wright, J. H. (2011). Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review*, *101*(4), 1514–1534.

Xiang, J., & Zhu, X. (2013). A regime-switching Nelson-Siegel term structure model and interest rate forecasts. *Journal of Financial Econometrics*, *11*(3), 522–555.

Yang, S., Liu, Z., & Wang, X. (2020). News sentiment, credit spreads, and information asymmetry. *The North American Journal of Economics and Finance*, *52*, 1–19.

Yu, W.-C., & Zivot, E. (2011). Forecasting the term structures of Treasury and corporate yields using dynamic Nelson-Siegel models. *International Journal of Forecasting*, *27*(2), 579–591.

Appendix

1.A Attributes and Summary Statistics for Sentiment Variables

	Interest Rates	**Inflation**	**Unemployment Rate**
Keywords	interest-rate, treasury-bill (-note, -bond) sovereign-debt, government-budget	inflation, deflation consumer-price-index	unemployment, employment jobless-claims, non-farm-payrolls
Number of News	22,548	7,821	8,281
Mean	49.97	43.94	65.52
Std. Deviation	22.58	22.50	25.48
Skewness	-0.41	0.13	-0.67
Kurtosis	-0.38	-0.41	-0.51
ACF(1)	0.71	0.44	0.46
ACF(2)	0.55	0.20	0.35

Table 1.9: Attributes and Summary Statistics for Sentiment Variables of RavenPack

The table shows the attributes and summary statistics for the selected sentiment variables from January 2008 - May 2022. Each column represents the sentiment on the underlying macroeconomic variables, i.e. interest rates, inflation rate, and unemployment rate. The first row shows the keywords used to extract the sentiment variables.

	Interest Rates	**Inflation**	**Unemployment Rate**
Keywords	interest rate, treasury bill	inflation	unemployment, labor market
Number of Tweets	1,019,957	1,710,748	1,632,518
Mean	-0.02	-0.02	-0.05
Std. Deviation	0.04	0.05	0.05
Skewness	-0.64	-0.94	-0.46
Kurtosis	1.48	3.45	1.03
ACF(1)	0.46	0.35	0.48
ACF(2)	0.43	0.33	0.46

Table 1.10: Attributes and Summary Statistics for Sentiment Variables of Twitter

The table shows the attributes and summary statistics for the selected sentiment variables from January 2008 - May 2022. Each column represents the sentiment on the underlying macroeconomic variables, i.e. interest rates, inflation rate, and unemployment rate. The first row shows the keywords used to extract the sentiment variables. In order to guarantee that the extracted tweets are about the economy of the United States of America, all tweets had to contain at least one of the following keywords: US, USA, United States, FOMC, FED, America, Powell, Yellen, Bernanke or Greenspan.

1.B RavenPack Sentiment vs. Macroeconomic Variables

Figure 1.3: Sentiment vs. Macroeconomic Variables (RavenPack)

This figure shows different macroeconomic variables in blue against their sentiment counterpart in yellow. All time series are standardized. They cover the period from January 2008 to May 2022.

	$S_{i,t}$	$S_{\pi,t}$	$S_{u,t}$	$S_{i,t+1}$	$S_{\pi,t+1}$	$S_{u,t+1}$	$S_{i,t-1}$	$S_{\pi,t-1}$	$S_{u,t-1}$
Δi_t	**-0.42**	-0.2	0.05	**-0.23**	-0.05	**0.51**	**-0.35**	0.07	-0.09
π_t	**0.22**	**-0.16**	**0.13**	**0.22**	-0.13	0.09	**0.2**	**-0.21**	0.13
u_t	0.15	-0.01	-0.2	0.12	-0.07	-0.07	**0.22**	0.09	**-0.18**

Table 1.11: Correlation Analysis Sentiment vs. Macroeconomic Variables

The table depicts the correlation between the RavenPack sentiment variables (contemporaneously, lead one period and lagged one period) and their macroeconomic counterparts. Δi_t, π_t and u_t represent the first difference in the short rate, inflation and unemployment growth, respectively. S_i, S_π, S_u denote the sentiment on each macroeconomic variable. Bold font indicates a p-value lower than 0.05, testing whether the true correlation coefficient is equal to zero. Due to the relatively small number of posts during the initial years of Twitter, the time period ranges from January 2011 to May 2022.

1.C Complete List of Models

- VAR DNS Only
- VAR Log Divisia
- VAR Percentage
- VAR Maturity
- VAR Log Assets
- VAR Michigan Cons. Sentiment
- VAR Fundsrate
- VAR All Sent. (Twitter)
- VAR Inf. Sent. (Twitter)
- VAR Int. Sent. (Twitter)
- VAR Unemp. Sent. (Twitter)
- VAR Int. Sent. (Ravenpack)
- VAR All Sent. (Ravenpack)
- VAR Inf. Sent. (Ravenpack)
- VAR Unemp. Sent. (Ravenpack)
- MSVAR DNS Only
- MSVAR Log Divisia
- MSVAR Percentage
- MSVAR Maturity
- MSVAR Log Assets
- MSVAR Michigan Cons. Sentiment
- MSVAR Fundsrate
- MSVAR All Sent. (Twitter)
- MSVAR Inf. Sent. (Twitter)
- MSVAR Int. Sent. (Twitter)
- MSVAR Unemp. Sent. (Twitter)
- MSVAR All Sent. (Ravenpack)
- MSVAR Inf. Sent. (Ravenpack)
- MSVAR Int. Sent. (Ravenpack)
- MSVAR Unemp. Sent. (Ravenpack)

1.D Performance of Raw Models

		Entire Trading Period		*Before Corona*		*During Corona*	
		VAR DNS Only	MSVAR DNS Only	VAR DNS Only	MSVAR DNS Only	VAR DNS Only	MSVAR DNS Only
Slope Bets	*Rank*	18	8	22	6	4	25
	T-Test p-value	0.15	0.49	0.06	0.67	0.85	0.29
Curvature Bets	*Rank*	18	14	7	11	24	14
	T-Test p-value	0.2	0.27	0.47	0.27	0.01	0.05

Table 1.12: Raw Model Performance
The table shows the performance of the two raw models *VAR DNS Only* and *MSVAR DNS Only* in terms of their ranks and the p-value of the test comparing the mean of the returns and the best-performing model in the category.

Chapter 2

Quantifying Artist Reputation: Disentangling Art-Market Subgroup Credentialing Effects

Revise and Resubmit at Poetics

Abstract

I create reputation time series data for a rich sample of emerging artists. The reputation measure reflects how much status-weighted attention the artist receives from leading art-market representatives. The main findings are threefold. First, art galleries act as gatekeepers. Second, other subgroups of the art market follow the lead of art galleries in terms of attention granted to an artist. Third, auction prices are influenced by the status-weighted attention granted pre-auction to the emerging artist by art galleries, art advisors, and the media in descending order of impact.

Keywords: auctions; art prices; artist reputation

2.1 Introduction

In markets for goods which lack objective quality standards, market intermediaries play a crucial role in shaping the perceived quality of the goods. This is especially relevant for cultural markets like the one for contemporary art. Bourdieu (1993, 1996) argues that in order to determine the economic value of a cultural product, we also have to consider the group of agents and institutions contributing to the value of the product through the creation of a belief in its quality. The market for contemporary art by emerging artists is a primary example of a market in which objective quality standards are difficult to determine. However, in the past, empirical studies have failed to show a significant effect of the work of art-market intermediaries on the economic value of the artists in terms of higher prices (Beckert & Rössel, 2013; Rengers & Velthuis, 2002; Velthuis, 2005).

The belief in an artwork's quality derives most importantly from the artist's credibility from the collector's perspective (Beckert & Rössel, 2013; Bonus & Ronte, 1997). An artist's credibility describes the confidence art collectors have in the artist that his or her works of art are of a certain artistic quality. An artist's reputation generates that credibility. The reputation of an artist is understood as a prediction that he or she can create works of art that will be acclaimed or exhibited by actors in the field with a certain status. Status refers to the hierarchical position of an actor in a social context. The artist's reputation is conferred by the status-weighted attention granted from the actors in the field of art (Beckert & Rössel, 2013; Bourdieu, 1996; Podolny, 1993).[1] Consequently, a measure of an emerging artist's reputation should ideally be based on the status-weighted attention he or she receives from relevant actors in the field. Attention should be status-weighted because the better an actor's hierarchical position within the field of art, the more weight should be given to the attention granted by that actor.

This paper presents a novel methodology for dynamically quantifying an artist's reputation over time. I measure the status-weighted attention an artist receives from various subgroups of art market intermediaries. I refer to this process as the credentialing process of artists. Specifically, I create a network of art-relevant institutions and individuals and apply centrality measures from graph theory to approximate the importance of the actors in the field of art. As a data source, I use Instagram. Tracking who, out of the network of art-relevant actors, has made posts related to an artist allows one to create artist-specific status-weighted attention time series that quantify an artist's reputation. The nodes of the network are divided into subgroups of art-market intermediaries. I define the subgroups at the level of art galleries, museums, media, curators, art advisors, collectors, and others. This allows the sources of reputation

[1] A field is to be understood as the plurality of relevant actors orienting their actions to one another (Beckert, 2020; Bourdieu, 1996; DiMaggio & Powell, 1983; Fligstein & McAdam, 2012). The field of art consists of the actors that are considered relevant in the art market for judging the quality of an artwork. It is important to note that not only art buyers (museums, private collectors) or sellers (artists, galleries, auction houses), but also intermediaries such as curators, art schools, art consultants, art historians, prize juries, art journalists or foundations belong to the field of art (Beckert, 2020). The aforementioned list of actors is not exhaustive.

to be disentangled at the subgroup level. In addition, I complement this novel methodology with a new and extensive dataset on auction results of emerging artists. The analysis is centered specifically on the sale of paintings as a means of facilitating a rigorous quantitative analysis. The decision to study the credentialing effects on prices for emerging artists is motivated by the following idea: For mature and well-known artists, inter-artist variation in received art-market attention is likely less pronounced in contrast to emerging artists. The latter makes it more difficult to econometrically identify credentialing effects on prices, which is why this study focuses on emerging artists.

Based upon the novel methodology to quantify an artist's reputation, this paper addresses three questions: First, which subgroup of the art market initiates an emerging artist's credentialing process? Second, what are the dynamics of artist credentialing between the different art market subgroups? Third, what is the relative influence of the subgroups' credentialing on the willingness to pay of contemporary art collectors for artworks of emerging fine art artists? To the author's best knowledge, the first question has not yet been answered quantitatively, and the last two questions have not been addressed in the literature due to methodological limitations.

First, I present empirical evidence that art galleries act as gatekeepers. They identify new emerging artists and initiate the credentialing process. This establishes the art gallery as the key belief-generating institution in the contemporary art market for emerging artists. Second, art galleries lead an artist's attention, i.e., the more status-weighted attention an artist receives from art galleries, the more status-weighted attention the artist receives from the other subgroups in the art market as well. An exception constitutes the museum subgroup. Museum attention for an artist is lagging and as such not directly influenced by an artists gallery attention. Third, in descending order, the impact of artist credentialing by art galleries, art advisors, and the media on prices is statistically and economically significant. I do not find any consistent effect on realized auction prices from credentialing by curators and museums. The relative effects of credentialing by the different subgroups on auction prices are rationalized with empirical evidence on the price-attention dynamics of the different subgroups and the interest structure of art collectors. The latter describes in which undertakings of the various art-market subgroups art collectors show how much interest.

Concerning the interest structure of art collectors, this paper shows for the first time that the interest structure of art collectors differs fundamentally from that of the representatives of the other art-market subgroups. I show quantitatively that, on average, representatives of the art-collector subgroup are most interested in the undertakings of art galleries, followed by undertakings of art advisors, media, curators, and museums. This is the same order as the relative effects of credentialing by the different subgroups on prices. Moreover, the interest structure of art collectors contrasts with that of other subgroups: for the other subgroups, museums receive on average the highest attention and art advisors the lowest. With respect to price-attention dynamics, I show that the museum and curator subgroups tend to flag commercially successful artists. They pay relatively more attention to artists whose prices are already high compared to the other subgroups. Using an adjusted Gini status-weighted attention inequality index, I also show that museums and curators distribute their status-weighted attention most unevenly among artists compared to the other subgroups. The results of both the interest structure and price-attention dynamics analyses are consistent with the results regarding the relative impact of credentialing by different subgroups on realized auction prices.

This paper proceeds as follows: Section 2.2 summarizes related literature. Section 2.3 describes the model of the credentialing process of emerging artists. Section 2.5 derives the construction of the artist reputation measure and specifies the different econometric models. Section 2.4 presents the data. Finally, Section 2.6 contains the empirical results of the models and their discussion. Section 2.7 concludes.

2.2 Literature Review

Typically, the social mechanisms that alter the value of cultural products like art, wine or fashion are grouped into processes of credentialing and qualification (Accominotti, 2022; Anand & Jones, 2008; Antal et al., 2015; Becker, 1982; Beljean et al., 2015; Bourdieu, 1993; English, 2005; Karpik, 2010; Khaire & Wadhwani, 2010; Lamont, 1987; Menger, 2014; Mulkay & Chaplin, 1982; Salganik et al., 2006; Smith, 1988). Henceforth, this review summarizes the literature relevant to the study of price formation of artworks by emerging contemporary artists. By credentialing at the level of the artist, I refer to processes in which intermediaries confer reputation on the artist when they endorse or give attention to the artist's œuvre or personality. By credentialing at the level of the artwork, I refer to processes in which intermediaries endorse a particular artwork by an artist. Since the intent of this paper is to study attention and price differences between artists, I focus on credentialing at the level of the artist; in what follows, the term credentialing refers to credentialing at the level of the artist unless otherwise stated. By credentialing, I explicitly do not refer to signaling processes between the art-market intermediaries.[2]

Qualification processes, on the other hand, occur when intermediaries shape the perceived value of things by establishing criteria for what constitutes quality (Accominotti, 2022; Beckert & Aspers, 2011; Boltanski & Thévenot, 2006; Stark, 2009). In this paper, qualification processes are not considered for two reasons. First, no economically significant effect of qualification efforts on contemporary art prices has been econometrically found so far. Velthuis (2005) econometrically tests the influence of the number of artists a gallery represents on their prices. The idea is that the fewer artists a gallery represents, the more time it can spend on each

[2] The possibility of signaling presupposes asymmetric information among market participants, which requires the existence of some objective quality standards. Since these have not been formally established, signaling processes are not included in the definition of credentialing.

artist to qualify his or her work by presenting it to the public. While Velthuis (2005) finds a statistically significant negative impact of the number of artists represented by a gallery, the impact is negligible in economic terms (p. 105). Furthermore, the qualification process, understood as shaping the standards for how an artist's work is perceived, can also be interpreted as a form of credentialing.

Given the sparse econometric support also for the credentialing theory, Accominotti (2022) argues that art collectors' differential willingness to pay for different contemporary artists is due to the existence of a reliable hierarchy of worthiness among the artists in the field. He argues that the relationships between artists and galleries create such a hierarchy in this field. This approach is limited in that it does not allow for analyzing whether other subgroups also influence artist price levels whose status-weighted attention patterns do not necessarily imply the same hierarchy among artists in the art field as that established by art galleries.

In order to study the determinants of art prices, financial art literature typically applies hedonic regressions of the price on a set of utility-bearing characteristics of the artworks, see for example Beckert and Rössel (2013) or Aubry et al. (2023).[3] These typically include physical characteristics of the artworks (e.g., size, medium, signature, date), and transaction characteristics (e.g., auction house, season, lot number). In addition, there is agreement within financial art literature that the artist's repu-

[3]This list is not exhaustive.

tation influences the price of artworks and should therefore be included in a hedonic regression. However, the literature lacks a methodology to capture this in a way that is consistent with the sociological notion of how an artist's reputation is created in the field of art.

In 1970, Willi Bongard began to measure the reputation of artists by developing a methodology for creating an artist ranking list. The ranking was based on several indicators: number of works included in permanent collections of museums or large private collectors, number of solo or group exhibitions, coverage in magazines, books and on television, and others. Bongard then weighted the achievements by assigning numerical values to them. The numerical weightings were determined by Bongard subjectively. For more details, refer to Grampp (1989).

Frey and Pommerehne (1989) tried to econometrically determine the parameters of an *aesthetic evaluation function.* The aesthetic evaluation was approximated with the Bongard points and considered to depend on the *artistic capital stock.* The constituents of the latter were defined very similarly to those of the Bongard measure. Therefore, by construction, Frey and Pommerehne (1989) explain much of the variation when Bongard scores are regressed on the artistic capital stock. However, their result cannot be interpreted in the sense that there exists an aesthetic evaluation function. Bonus and Ronte (1997) then claimed that there are no objective criteria for determining the quality of artworks. They further argued that due to the lack of objective criteria, the public credibility of an artist determines the economic value of his or her work. Credibility is created by experts in the art community. Due to the apparent randomness in whether, to what extent, and for what period of time an artist is recognized, Baumol (1986) referred to the art market a floating crap game from an investment perspective.

Beckert and Rössel (2013)[4] derived sociologically that an artist's artistic reputation is socially constructed over time through the conferring of reputation by experts and institutions in the field of art. They tested econometrically the influence of the reputation conferred by galleries, universities, media, and awards on auction prices. They found a statistically significant and positive effect of media awareness and awards on auction prices.[5]

In addition to the literature summarized above, much of the quantitative financial art literature treats artist reputation in a static way, approximating it simply through artist fixed effects (Aubry et al., 2023; Chanel & Gérard-Varet, 1996; de la Barre et al., 1994; Etro & Pagani, 2013; Ma et al., 2022; Marinelli & Palombo, 2011; Renneboog & Spaenjers, 2011; Sproule & Valsan, 2006). Artist fixed effects, however, capture all uncontrolled price-influencing factors that are constant and artist dependent. Therefore, artist fixed effects reflect more than just the rep-

[4]In 2013 an English translation of the original German version of the paper appeared in (Beckert & Rössel, 2004).

[5]Beckert and Rössel (2004) also find a negative, statistically significant, effect of the reputation variable *solo exhibition*. Beckert and Rössel (2013) do not include *solo exhibition* due to multicollinearity problems.

utation of the artist. Moreover, since almost all data samples contain multiple artworks by an artist from different sales, and reputation is not constant, artist fixed effects at best reflect average reputation and other personal characteristics. Other contributions to the financial art literature use semi-static proxies for artist reputation by creating dummy variables for selected reputation indicators. The literature has considered reputation-approximating indicator variables like art prizes (Taylor & Coleman, 2011), exhibitions at well-known art shows (Renneboog & Spaenjers, 2013), or whether the artist's work is referenced in art books (Campos & Barbosa, 2009; Renneboog & Spaenjers, 2013). A commonality of these reputation-approximating factors is that none of them can decrease. However, an artist's reputation can also diminish over time. An artist needs continuous reconfirmation from the field of art (Beckert, 2020).

Fraiberger et al. (2018) capture the movement of certain artists' art through a network of galleries and museums. They show that artists with early access to high-status institutions tend to retain access to these institutions, command higher prices, and have lower exit rates than those who begin on the periphery of this network, whose access to prestigious venues often remains restricted in perpetuity. Although they argue that artists' success in the arts ecosystem is controlled by gatekeepers, they do not disentangle the reputational effect of different art-market intermediaries and institutions on career trajectories and prices.

The literature on the role of gatekeepers in the cultural industry is abundant. In a qualitative case study of abstract expressionists, Bystryn (1978) shows that representation by a high-status gallery often entails other forms of success, such as attention by reputable critics and acquisition by prestigious museums and collectors. For the digital crypto-art sector, Vasan et al. (2022) note that the relationships between artists and collectors are critical to artists' success, since the traditional art-market intermediaries such as galleries, museums, and curators are absent from this market. Other notable studies concerning the role of gatekeepers in the cultural industry include: Coslor et al. (2019); Foster et al. (2011); Godart et al. (2023); Kawashima (1999); Smits (2016).

2.3 Credentialing Process Model

The model of the credentialing process is shown in Figure 2.1. The field of art consists of various subgroups. Each subgroup consists of different actors. The bestowal of reputation is understood as the status-weighted attention an artist receives from actors in the field of art. The conferring of reputation could take the form of a gallery organizing an exhibition for an artist, a critic writing about an artist in an art magazine, an art consultant advising a client to buy an artwork by a particular artist, etc. Status-weighted attention from the field of art creates an artist's reputation; the attention can be broken down by the subgroups that bestowed it. The totality of reputation conferred by all subgroups makes up an artist's reputation. The reputation of an artist is understood as a prediction that he or she can create artworks that will be acclaimed or exhibited by actors in the field with a certain status. An analogous definition applies to reputations at the subgroup level.[6] A collector is assumed to transform the observed status-weighted attention an artist receives into his or her personal assessment of the artist's credibility. A collector reweights the status-weighted attention an artist receives from a subgroup according to his or her own subgroup-specific impact factor. The impact factor reflects the extent to which a collector's credibility assessment is influenced by the status-weighted attention of a particular art market subgroup. An artist's credibility assessed in this way determines a collector's belief in the overall artistic quality of an artist. This drives an art collector's reputation-based willingness to pay. Artwork-specific characteristics such as size, material, age, etc. determine an art collector's ultimate willingness to pay for a particular artwork by an artist.

[6]The model up to this point is adapted from Beckert (2020); Beckert and Rössel (2013); Bourdieu (1993). A related model on the development of status is introduced by Gould (2002).

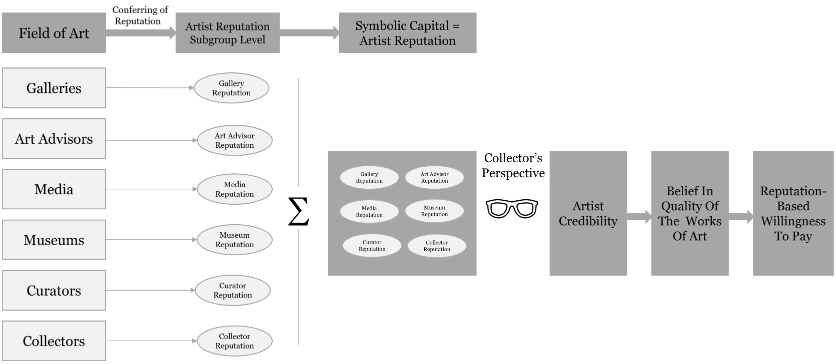

Figure 2.1: Model of the Credentialing Process

2.4 Data

There is no formal definition of the term emerging artist. Therefore, in order to define the group of emerging contemporary artists, I selected all artists who have sold at least one object in the auction series *New Now* of Philipps auction house, a semi-annual auction in New York City and London dedicated to the sale of artworks by emerging contemporary artists. This selection rule implies that all artists have successfully sold at least one lot in an internationally renowned auction house. From that point of view, rather advanced emerging artists are considered in this paper.[7] There were a total of 21 *New Now* auctions between 2015 and

[7]The findings presented in this paper were derived from a comprehensive collection of emerging artists. The external validity of the results of this paper for alternative definitions of emerging artists remains a topic for future investigation. However, given the vast array of artists examined and given that robustness tests with subsamples of the artists yielded qualitatively equivalent results as shown in Appendix 2.E, no significant alterations of the results are expected.

2021. Since this paper analyzes paintings only, artists working exclusively in other mediums (sculpture, photography, etc.) were excluded. In addition, due to data procurement constraints, the 90^{th} percentile in terms of total number of hashtag posts on Instagram for a single artist on Instagram was used to cut the artist sample from the top.[8] This procedure also ensured that artists that are already too famous and cannot be considered emerging artist anymore were dropped from the sample. Furthermore, only living artists were considered. This resulted in a sample of 457 artists. For these artists, the study analyzed all lots sold at all major auction houses (not just Philipps). The data were downloaded from *LiveArt*. In total, there are $6,739$ sold lots by these 457 artists for 2015-2021. Table 2.1 contains descriptive statistics.

[8] In a robustness check accounting for this decision, the results of this paper were regenerated using a random 90% subsample of the final set of artists. The results stayed qualitatively the same. The results are tabulated in Appendix 2.E.

	Total
Number of observations	6,739
Number of distinct artists	457
Number of distinct auction houses	113
Number of auction events	694
Share sold at Sotheby's	23.3%
Share sold at Christie's	28.2%
Share sold at Phillips	34.3%
Median hammer price ($)	31,077
Mean hammer price ($)	132,587
Maximum hammer price ($)	16,767,500
Minimum hammer price ($)	134
Share of Female artists	18.6%
Share of USA artists	50.4%
Share of UK artists	9%
Share of German artists	8.6%
Share of French artists	3.2%
Share of Swiss artists	2%
Share sold in January	1.4%
Share sold in February	6.8%
Share sold in March	12.6%
Share sold in April	7.6%
Share sold in May	10%
Share sold in June	10.3%
Share sold in July	6.3%
Share sold in August	0.7%
Share sold in September	10.9%
Share sold in October	11.6%
Share sold in November	12.6%
Share sold in December	9.2%

Table 2.1: Descriptive Statistics of Artworks

This table shows descriptive statistics of the data used in this paper. The prices are given in real 2017 USD.

2.4.1 Mapping the Global Art Community

Instagram has quickly become the most used social media platform in the field of art since its inception in 2010.[9] The platform allows sharing of photos and videos, among other things. Individuals and institutions create personal accounts on which they curate their content. One account can follow other accounts, while the other account does not need to follow back; popular accounts with many followers often follow far fewer other accounts. The starting point for the digital mapping of the global art community network in this paper is 10,000 generally art-relevant Instagram accounts. These museums, galleries, curators, art advisors, media, off-spaces and others.[10] The Instagram accounts followed by the initial 10,000 generally art-relevant accounts, may be assumed to be of importance in the art world as well. Therefore, all accounts followed by the initial 10,000 accounts have also been included.[11] This procedure yields a total of 2,857,254 unique accounts. The identified account-follower relationships are used to create a network, where a node represents an identified Instagram account. I use indegree and eigenvector centrality measures from network-theory to approximate the importance of the nodes in the network, which is a common method in network theory

[9] Hiscox online art trade report 2019: https://www.hiscox.co.uk/sites/uk/files/documents/2019-04/hiscox-online-art-trade-report-04-2019.pdf

[10] The initial account names were provided by NAMAC GmbH., a consulting firm for art market trends. The accounts were assembled by NAMAC in spring 2021.

[11] The data was downloaded using proprietary software in spring 2021. The data were downloaded all at once. While followingships change over time, they can be safely assumed to be more or less stable over the short-term as applicable in this study. In other words, it is unlikely that the art world representatives that have the power to credential artists changed markedly over the 6 years analysed in this paper.

(Jackson, 2011).[12] Eigenvector and indegree centrality are amongst others different ways to measure the centrality of a node in a network, which is used to approximate the status of an art-market player in the network. The econometric results of the analysis do not depend on the choice of the type of the centrality measure. For a more formal treatment of network theory, I refer to Jackson (2011).

The top 4,095 nodes in terms of their network rank with respect to indegree centrality were manually classified into the art-market subgroups art galleries, art advisors, media, museums, curators, and others. While the relative centrality of a node may increase and decrease over time, centrality values were assumed to be constant due to the relatively short observation period of the empirical analysis in this paper. The top 100 Instagram accounts in terms of their indegree centrality statistics are listed in Appendix 2.A.

Table 2.2 provides descriptive statistics on the top 4,095 Instagram accounts reflecting the institutions and individuals in the different art-market subgroups. The classification into these subgroups followed an inductive procedure, meaning that the subgroups were created based on observations. Public information was used to assign individual Instagram accounts to art-market subgroups.

[12]For computational purposes, the network had to be pruned. The network was pruned by keeping all nodes above the 90% percentile in terms of their indegree centrality statistic. The final network consists of 221,357 nodes.

Many art collectors do not publicly disclose their collecting activity. Therefore, I also used a registry of contemporary art collectors provided by Larry's List to identify Instagram accounts of contemporary art collectors. Since information about the buyer of a lot at an art auction is generally not public, the exact proportion of lots sold to art collectors in the dataset cannot be determined. The 337 art collectors collectively created 27,159 posts with hashtag references to the artists in the dataset. Under the reasonable assumption that an art collector or art-collection representatives create Instagram posts about, among other things, the artists they collect, it can be further assumed that a portion of the art was sold to these collectors.

20 accounts are included in both the curator and art-advisor subgroups, because both occupations are reported online for these accounts, and it was not possible to determine whether one occupation dominated the other in terms of time spent or prestige. In addition, only museum directors who are not working as curators were included in the museum subgroup. Museum directors who are still working as senior curators in that museum were instead included in the curators group. Finally, media such as art magazines and newspapers work with other art-market professionals such as art critics, art historians, etc. From this perspective, the attention of these art-market experts is indirectly transferred through the media. This has to be taken into account when interpreting the econometric results in Section 2.6.

Column 4 in Table 2.2 shows how many Instagram posts each subgroup made for one of the 457 artists in the data sample.[13]

Art-Market Subgroup	*Number of Observations*	*Type of Accounts*	*Number of Artist Posts*
Galleries	1,020	Official Gallery Accounts, Gallery Owners	64,617
Curators	424	Known Professional Curators	19,632
Museums	390	Official Museum Accounts (Private and Public), Museum Directors	10,361
Collectors	337	Individual Art Collectors Official Institutional Collection Accounts	27,159
Media	262	Accounts of Art Magazines (Online and Offline) Newspapers, Journalists, TV Channels	13,477
Art Advisors	212	Art Advisors Art Advisories (Institutional Accounts)	25,934
Others	1,450	Offspaces, Art Foundations, Non-Profit Organisations, Artists Art Professors, Art Historians, Art Fairs	55,228
Total	4,095		216,408

Table 2.2: Art-Market Subgroup Descriptions

This table shows the different subgroups of art-market players and selected descriptive statistics of their Instagram profiles.

Figure 2.2 illustrates the distributions of centrality scores of actors and institutions in different art-market subgroups. The subgroup order in terms of median centrality value significantly changes depending on the centrality measure. While art galleries have the lowest median eigenvector centrality, they have the highest median indegree centrality. The subgroup order in terms of the relative range of centrality values also changes depending on the centrality measure.

[13] For additional descriptive statistics on the Instagram profiles of the 4,095 art market Instagram accounts, see Appendix 2.H

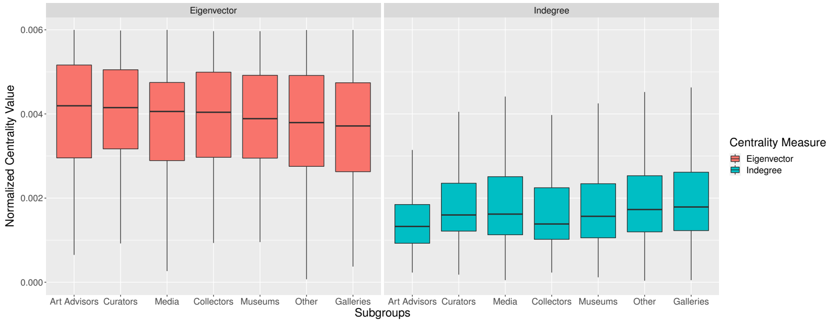

Figure 2.2: Distributions of Centrality Measures by Subgroup

The distributions of two types of centrality values for actors and institutions across the art-market subgroups are shown.

Finally, the methodological consequences of using Instagram as data source warrant a discussion. The question of whether or not my Instagram-based network representation of the field of art accurately reflects the real structure of the art market cannot be answered definitely. However, the accounts appearing on the top of the ranking constitute no surprises, which would contradict with art-market intuition. However, it must be noted that comparing the status of a gallery and a curator may not be meaningful.[14] Although the centrality scores of network actors may suggest making status comparisons across subgroups, the analysis in the paper requires maintaining consistent relative rankings within subgroups rather than determining absolute rankings across all subgroups. Therefore, the crucial aspect is the relative ordering within each subgroup, and the proposed methodology successfully accomplishes this goal.

The author does not claim or imply that an artist's reputation is actually

[14] I thank an anonymous reviewer for pointing this out to me.

derived from his or her Instagram presence. It is true that younger artists often use social media platforms like Instagram for marketing and publicity purposes, but ultimately an artist's reputation is generated through the status of the institutions and people who grant the artist attention and thereby reduce the artistic quality uncertainty (Beckert & Rössel, 2013). Instagram data allows to approximate such attention.

It is an interesting avenue for future research to investigate whether the sheer number of likes or followers an artist has on social media actually translates into increased reputation. However, the issue of artificially inflating interaction numbers with fake accounts, which may be steered by the artists themselves, poses a serious threat to the concept that such numbers accurately reflect an artist's quality and reduce the artistic quality uncertainty surrounding their work.

2.5 Methodology

2.5.1 Measuring Status-Weighted Artist Attention

The social media platform Instagram is used as a data source. On Instagram, people can upload pictures to their personal profile and add a description. A common way to write the description is to use the hashtag symbol #. For example, someone might go to an art gallery to look at a painting by Rashid Johnson and take a photo of it, then upload it to Instagram with the hashtag #rashidjohnson. There is a feature on Instagram to search for hashtag posts. Searching for #rashidjohnson will return all posts that have used the hashtag #rashidjohnson.

In order to approximate the status-weighted attention an artist receives, I use all Instagram posts with the hashtag created by writing together the first and last name of that artist. Requiring the hashtag to include both the first and last name reduces the number of results, but it avoids false results if an artist has a well-known last name.[15] For every post, the data includes the date of when a post was uploaded and by whom identified with the username. The value of a post is equal to the node weight, determined by network centrality, of the user who uploaded the post. If the username of the poster is not included in the art network community, the value of the post is zero. Equation (2.1) describes the monthly attention value of artist i in month t by art-market subgroup k for $k \in \{\text{Galleries, Art Advisors, Media, Museums, Curators, Others}\}$. Let $N_{i,t,k}$ be the set of actors belonging to subgroup k that paid attention to artist i in month t. The status-weighted attention of artist i in month t by subgroup k is calculated as:

$$A_{i,t,k} = \sum_{j \in N_{i,t,k}} \alpha_j, \tag{2.1}$$

where α_j represents the centrality node weight of institution (person) j in subgroup k. These attention scores per art-market subgroup k can be interpreted as a continuous hierarchy representation per subgroup. The proposed set up allows to account for multiple hierarchies at the

[15]The data was downloaded in spring 2021.

same time. From this perspective, this paper generalizes the setup of Accominotti (2022).

The total status-weighted attention of all artists by the top 4,095 art-market actors and institutions in the dataset over time is shown in Figure 2.3a. In addition to the sharp increase in overall activity levels through 2017, strong seasonal patterns are present. The lowest activity level of the art community on Instagram is always in August, which coincides with the month with the lowest proportion of lots sold annually in my data, as shown in Table 2.1. Individual attention scores are adjusted for changing Instagram activity levels over the years. I create subgroup-specific Instagram activity indices using 2017 as the base year. All attention scores have been adjusted by multiplying them by the inverse of each subgroup's annual activity index. The adjusted total attention sum per month over all artists is shown in Figure 2.3b.

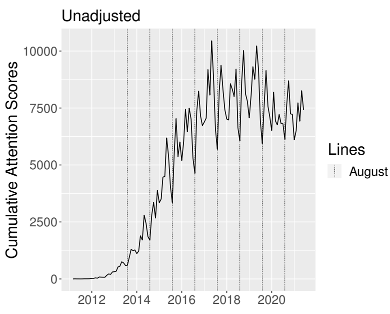

(a) Unadjusted

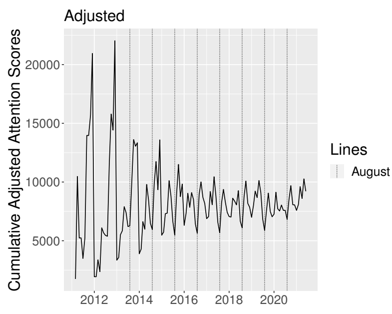

(b) Adjusted

Figure 2.3: Total attention sum over all artists analysed.

Due to the erratic behavior of attention levels through mid-2014, the time period for the analysis is set to 2015-2021.[16] In order to account for the remaining seasonality in the data, attention scores are aggregated on an annual basis, and month and year fixed effects are included in the models.

2.5.2 Models

Attention Dynamics

Let $A_{i,k,T}$ be the average attention that artist i receives from the art market subgroup k in year T:

$$A_{i,k,T} = \sum_{t=1}^{12} \frac{1}{12} A_{i,t,k}. \tag{2.2}$$

I use OLS to estimate the following regressions including 1-year lagged mean prices to determine the attention dynamics between the different art market subgroups:

$$A_{i,k,T} = \alpha + \sum_{j=1}^{2} \sum_{k=1}^{K} \beta_{j,k}\, A_{i,k,T-j} + \beta_m \,\text{Mean Price per SQCM}_{i,T-1} + \epsilon_{i,T}. \tag{2.3}$$

In words, equation (2.3) regresses year 3 attention of an art-market subgroup k on the attention measures of the other subgroups from year 2 and 1 and lagged mean prices. The attention dynamics are evaluated using yearly attention averages per subgroup because of the seasonality

[16]The attention dynamics analysis also makes use of data from 2014 and 2013 as explained in Section 2.5.2.

in art-market attention over the year. Furthermore, the analysis is restricted to 2-year attention lags for two reasons. Firstly, it is essential to realize that Instagram was only established in 2010. Consequently, the attention variables for the years preceding 2015 are built with comparably little information, necessitating considerable extrapolation, which in turn generates data noise in the attention variables. As a result, incorporating additional attention lags in the regression framework appears suboptimal from this perspective. Moreover, although it is theoretically plausible that the attention received by an art-market subgroup three or more years prior may still exert an influence on current art market attention, the likelihood of this occurrence is markedly low. When combined with the apprehensions regarding data quality for higher lags, it was opted to include a maximum of two years of lagged variables in the analysis. Consequently, the earliest attention scores used in the attention dynamic evaluations are from 2013.

Price Levels of Artworks

I use hedonic pricing models to study the relative influence of the reputation generated by the different subgroups of the art market on realized auction prices. Reputation is quantified using the status-weighted attention scores derived in equation (2.1). The dependent variable is the natural logarithm of realized auction hammer prices in real (2017) US Dollars per square centimeter. The hedonic model is:

$$\ln(p_{j,t}) = \alpha + \sum_{k=1}^{K} \beta_k \sum_{\tau=1}^{12} \delta_\tau \mathrm{A}_{i,k,t-\tau} + \sum_{m=1}^{M} \beta_m \mathrm{Artist}_{m,j,t} + \sum_{n=1}^{N} \beta_n \mathrm{Transaction}_{n,j,t} + \sum_{p=1}^{P} \beta_p \mathrm{Physical}_{p,j,t} + \epsilon_{j,t}, \tag{2.4}$$

where $ln(p_{j,t})$ represents the natural logarithm of the hammer price per square centimeter of artwork j sold in month t. The main independent variables of interest are the 12-month δ_τ discounted status-weighted attention scores of artist i by subgroup k before month t. The coefficients are expected to be non-negative, reflecting the notion that any form of publicity is beneficial. The model is estimated using both unit discount factors and data-driven discount factors. The derivation of the data-driven discount factors is described in Appendix 2.B. By including only the attention scores for the months preceding the auction month, reverse causality issues are avoided. I estimate the model twice; attention values $A_{i,t,k}$ are calculated using either eigenvector or indegree centrality measures to determine node weights. The different β_k describe the effects of subgroup k's credentialing on an artist's price per square centimeter.

The β_k of the collector subgroup from model (2.4) reflects both the effect of credentialing and the effect of actual collector demand on prices per square centimeter. Representatives in the other subgroups may also buy art at an auction, but in comparison to the art-collector subgroup, they are far less associated with actual buying activity in the secondary market. Consequently, the β_k of the other subgroups are interpreted as describing only the effect of credentialing on prices.

In addition, I control for $\text{Artist}_{m,j,t}$, which are artist-specific attributes m of artwork j at time t; $\text{Transaction}_{n,j,t}$, which are transaction-specific attributes n of artwork j at time t; and $\text{Physical}_{p,j,t}$, which are physical attributes p of artwork j at time t.

The regressors are summarized below:

- **Attention**.
  - $A12_k$ for $k \in$ (Galleries, Art Advisors, Media, Museums, Curators, Others) describes the sum of discounted attention scores of art-market subgroup k for 12 months prior to the auction month.

The control variables are specified as follows:

- **Artist**
  - *Age*
  - *Gender*
  - *Price per square centimeter T-1 Mean:* The mean realized price per square centimeter in the previous year[17]
- **Transaction**

[17] In Appendix 2.G the price regressions are re-estimated using lagged median or max prices. The results concerning the credentialing effects remain qualitatively untouched.

  - *Month dummies:* Month fixed effects to control for art-market seasonalities
  - *Auction house dummies:* Fixed effects for the three largest auction houses - Sotheby's, Christie's, and Phillips (the other auction houses serve as a benchmark)
  - *Liquidity:* Counts the number of paintings sold by that artist in the past month

- **Physical**
  - *Size:* The width and length in centimeters of each artwork, and their quadratic terms
  - *Artwork materials:* 5 indicator variables, used for the 5 most common material descriptions: oil, acrylic, ink, gouache, and mixed media (not mutually exclusive)
  - *Age of artwork:* The age of the artwork at the time of sale

Attention dynamics among art-market subgroups, described in model 2.3, and the relative effects of different subgroups' credentialing on prices, model 2.4, are estimated using subgroup representatives consisting of the top 4,095 individuals and institutions in the art market in terms of their network ranking. Inclusion of subgroup representatives below this threshold could, in principle, affect the estimates. However, it is unlikely to change the main results. This is because of the already large sample of 4,095 subgroup representatives, from which the most relevant in terms of network centrality were selected. Moreover, the relative importance of subgroup representatives in terms of their network centrality decreases rapidly and exponentially, as shown in Appendix 2.C.

The set of control variables does not include other reputation-capturing variables such as scholarships or the number of awards won. Beckert and Rössel (2013) control for such covariates. However, this is not an issue, because the purpose of the study is to disentangle the credentialing effects of different art-market subgroups on art prices. Awards or scholarships do influence prices via the credentialing channel.

2.6 Results

2.6.1 Attention Dynamics

The results of the attention dynamics model (2.3) are shown in Table 2.3. Unless otherwise noted, attention in the following represents status-weighted attention. Gallery attention to an artist evolves almost independently of the attention of the other subgroups of the art market. An artist's 1-year lagged gallery attention strongly predicts current gallery attention. A one-standard-deviation increase in 1-year lagged gallery attention is associated, ceteris paribus, with an average increase in current gallery attention of about 0.56 standard deviations.[18] This is intuitive, as galleries usually have long-term business relationships with their artists; consequently, earlier attention should be strongly correlated

[18] For interpretation, I use the arithmetic mean of the two coefficients due to the different network centrality measures eigenvector and indegree.

with current attention. Moreover, no other subgroup's lagged attention can explain current artist attention of art galleries. While the 1-year lagged attention of the collector subgroup and the outsider subgroup, as well as the 2-year lagged curator and collector attention statistically significantly correlate with current art gallery attention for an artist, the magnitude of the respective coefficients is economically negligible and thus defy interpretation. Moreover, the lagged attention of the art galleries strongly predicts the current attention of all other subgroups except museum attention. In other words, if an artist has experienced an increase in status-weighted attention from art galleries in the past year, the other subgroups can be expected to increase their status-weighted attention to that artist on average in the current year. In summary, the lagged attention of galleries strongly predicts the attention of all other subgroups (except museums), while the lagged attention of the other subgroups (except museums) do not explain the variation in current gallery attention. This supports the idea that art galleries initiate and lead an artist's credentialing process, which is consistent with the activities that art galleries typically undertake in the primary art market. An art gallery is expected to build an artist's career by organizing art exhibitions, performing various marketing tasks, and building relationships with art collectors. Thus, the art gallery as an institution can be considered the key belief generator in the contemporary art market. In particular, the attention of the media and art advisors is the plaything of art galleries, in the sense that art advisors' and media's 1-year lagged attention is less predictive of their current attention than is 1-year lagged attention of art galleries. Finally, the finding that lagged gallery attention is not predictive for current museum attention for an artist, is consistent with the idea that museums grant attention to an artist only once an artist's reputation is rather firmly established. A result also consistent with the findings concerning the role of museums in the price formation process of emerging artists discussed in Section 2.6.2.

In general, the lagged attention of a subgroup correlates statistically and economically significantly with its current attention. This finding is the least pronounced for the media subgroup. This can be explained by the fact that individuals and institutions belonging to the media subgroup are the least likely to have a personal or contractual relationship with an artist, which is likely to cause attention persistence in the other subgroups.

At the individual level, lagged art-advisor attention does not explain statistically significantly the current attention of any other subgroup. The same holds true for the media subgroup. This finding suggests that both the media and art advisor subgroup are not taste makers in the art market. They react in the sense that they pay statistically significantly more attention to artists who were paid more attention to particularly by art galleries, in the previous year.

1-year lagged museum attention statistically and economically significantly predicts media attention. This is intuitive since a big museum show of an artist may very well attract media attention. However, the size of the effect is still much smaller in comparison to how much an increase in gallery attention sparks the other subgroups' attention. Furthermore, 2-year lagged museum attention statistically significantly predicts current curator attention. However, the measured effect is very small and is thus negligible. Interestingly, lagged attention from museums does not predict attention from art collectors. This result is consistent with the results of the hedonic price regressions in Table 2.4. Finally, museum attention for artists, similar to gallery attention, develops autonomously from the attention of other subgroups. Furthermore, it has been observed that museums, in general, do not amplify the attention received by other subgroups, with the exception of the media. These results suggest that museums do not follow the trends and hype of the emerging artist art market, and they choose seemingly independently which artists to showcase and grant attention.

Finally, 2-year lagged curator attention is statistically significantly and positively correlated with the current attention of all other subgroups. However, except for the effect of 2-year lagged attention on the current media attention, the effects are effectively small. The result suggests that (long) lagged curator attention can predict current art-market attention, although the signal is weak in comparison to how much the art gallery subgroup shapes the current art market attention. It must be noted, that the 424 curators in the sample can in fact exercise very different functions. For example, there are known professional museum curators in the sample as well as independent curators. While museum curators are less associated with the belief generation and taste making in the contemporary art market, independent curators can often fulfill these roles with significant impact. The finding that long lagged curator attention is somewhat predictive for current art-market attention suggests that there are at least some taste maker curators within the sample that can impact the attention trail over the long run. An interpretation consistent with these results is that some curators might discover and promote artists, who then get pushed into the wider public attention through gallery credentialing.[19,20,21]

The 1-year lagged attention of collectors correlates positively and statistically significantly with the attention of art advisors and the media. This result suggests that lagged collector attention can influence the current attention of those subgroups. However, the correlations are relatively small in economic terms and in comparison to how much the gallery subgroup influences current art-market attention. Moreover, concerning the 2-year lagged collector attention effects on current art-market attention, I find negative correlations. However, the size of the effects are economically negligible. Finally, mean prices per square centimeter

[19]This discussion indicates that future investigations may benefit from establishing multiple curator categories in order to discern the effects at a more granular level.

[20]1-year lagged curator attention is statistically significantly and negatively correlated with current collector and art advisor attention. However, the measured effects are very small in economic terms.

[21]Furthermore, Ullrich (2017) advances the idea that commercially successful artists might be denied the ability to empathize with minorities or underprivileged milieus, and thus to make art that meets the sociopolitical demands of curators. A likely consequence would be that commercially successful artists receive less status-weighted curator attention compared to the less commercially successful artists. Figure 2.9 in the Appendix 2.D contradicts this notion.

realized one year in advance only correlate statistically significantly with the current attention of art advisors and the coefficients are small in economic terms. A one-standard-deviation increase in lagged prices per square centimeter is associated with a 0.03 standard-deviation increase in the current attention of art advisors.

	Gallery Attention		*Museum Attention*		*Collector Attention*		*Media Attention*		*Art Advisor Attention*		*Curator Attention*	
	Eigenvector	Indegree	Eigenvector	Indegree	Eigenvector	Indegree	Eigenvector	Indegree	Eigenvector	Indegree	Eigenvector	Indegree
Gallery -1	0.540^{***}	0.589^{***}	0.123	0.212	0.542^{***}	0.607^{***}	0.424^{*}	0.475^{*}	0.430^{***}	0.551^{***}	0.577^{***}	0.539^{***}
	(0.035)	(0.036)	(0.110)	(0.146)	(0.176)	(0.178)	(0.239)	(0.269)	(0.133)	(0.151)	(0.205)	(0.196)
Art Advisor -1	0.0001	−0.003	0.012	0.022	0.070	0.074	0.069	0.090	0.297^{***}	0.373^{***}	−0.029	−0.042
	(0.008)	(0.008)	(0.025)	(0.033)	(0.046)	(0.046)	(0.059)	(0.066)	(0.036)	(0.041)	(0.041)	(0.038)
Media -1	−0.002	−0.004	0.015	0.011	0.018	0.007	0.136^{***}	0.152^{***}	0.001	−0.015	−0.018	−0.023
	(0.005)	(0.005)	(0.017)	(0.023)	(0.028)	(0.027)	(0.041)	(0.046)	(0.022)	(0.024)	(0.031)	(0.030)
Museum -1	0.008	0.008	0.247^{***}	0.317^{***}	−0.009	−0.027	0.123^{**}	0.121^{*}	0.046	0.049	0.061	0.035
	(0.008)	(0.008)	(0.030)	(0.040)	(0.040)	(0.039)	(0.060)	(0.066)	(0.033)	(0.037)	(0.047)	(0.044)
Curator -1	−0.005	−0.005	0.027	0.044	-0.078^{**}	-0.072^{*}	0.037	0.035	-0.048^{*}	-0.057^{*}	0.427^{***}	0.444^{***}
	(0.008)	(0.008)	(0.027)	(0.035)	(0.038)	(0.037)	(0.051)	(0.056)	(0.029)	(0.032)	(0.061)	(0.062)
Collector -1	0.017^{***}	0.014^{**}	0.034	0.036	0.575^{***}	0.601^{***}	0.149^{***}	0.159^{***}	0.121^{***}	0.143^{***}	0.054	0.040
	(0.006)	(0.006)	(0.022)	(0.029)	(0.040)	(0.041)	(0.047)	(0.052)	(0.027)	(0.032)	(0.037)	(0.035)
Other -1	0.028^{***}	0.023^{**}	0.006	−0.003	0.064	0.029	0.075	0.084	0.057	0.046	0.062	0.059
	(0.009)	(0.009)	(0.026)	(0.034)	(0.053)	(0.053)	(0.058)	(0.065)	(0.037)	(0.042)	(0.053)	(0.049)
Gallery -2	0.131^{***}	0.162^{***}	0.074	0.117	-0.540^{***}	-0.490^{***}	−0.084	−0.059	-0.225^{*}	-0.241^{*}	−0.257	−0.207
	(0.034)	(0.035)	(0.090)	(0.120)	(0.166)	(0.165)	(0.211)	(0.234)	(0.124)	(0.143)	(0.176)	(0.166)
Art Advisor -2	0.013^{*}	0.012	0.002	−0.005	0.046	0.060	0.023	0.015	0.133^{***}	0.155^{***}	−0.013	−0.010
	(0.007)	(0.007)	(0.025)	(0.034)	(0.042)	(0.042)	(0.052)	(0.059)	(0.033)	(0.037)	(0.041)	(0.038)
Media -2	0.003	0.003	0.016	0.015	−0.010	−0.024	0.133^{***}	0.146^{***}	0.019	0.012	0.012	−0.001
	(0.005)	(0.005)	(0.014)	(0.019)	(0.029)	(0.028)	(0.036)	(0.041)	(0.023)	(0.026)	(0.031)	(0.029)
Museum -2	0.002	0.003	0.148^{***}	0.190^{***}	0.015	0.024	0.029	0.030	0.038	0.048	0.073^{**}	0.067^{**}
	(0.007)	(0.007)	(0.039)	(0.054)	(0.035)	(0.034)	(0.050)	(0.055)	(0.028)	(0.032)	(0.036)	(0.034)
Curator -2	0.017^{**}	0.019^{***}	0.052^{**}	0.067^{**}	0.077^{**}	0.075^{**}	0.101^{*}	0.109^{**}	0.061^{**}	0.064^{**}	0.263^{***}	0.281^{***}
	(0.007)	(0.007)	(0.023)	(0.030)	(0.035)	(0.035)	(0.051)	(0.056)	(0.026)	(0.029)	(0.054)	(0.057)
Collector -2	-0.018^{***}	-0.016^{***}	-0.050^{***}	-0.061^{***}	0.024	0.046	-0.130^{***}	-0.141^{***}	-0.053^{**}	-0.051^{**}	-0.093^{***}	-0.088^{***}
	(0.005)	(0.005)	(0.016)	(0.021)	(0.032)	(0.033)	(0.034)	(0.038)	(0.021)	(0.024)	(0.025)	(0.025)
Other -2	0.014^{*}	−0.001	−0.005	−0.011	0.116^{**}	0.081	0.106^{*}	0.100	0.062^{*}	0.051	0.070	0.040
	(0.008)	(0.008)	(0.026)	(0.034)	(0.051)	(0.050)	(0.059)	(0.065)	(0.033)	(0.038)	(0.061)	(0.060)
Price SQCM -1 Mean	0.001	0.003	−0.004	−0.006	0.015	0.004	0.001	0.00001	0.029^{**}	0.031^{**}	0.014	0.014
	(0.003)	(0.003)	(0.008)	(0.011)	(0.020)	(0.020)	(0.019)	(0.020)	(0.012)	(0.014)	(0.016)	(0.014)
Constant	-0.010^{***}	-0.005^{***}	-0.019^{***}	-0.014^{**}	-0.039^{***}	-0.030^{***}	0.014	0.036^{***}	-0.030^{***}	−0.010	-0.021^{**}	-0.016^{*}
	(0.002)	(0.002)	(0.005)	(0.006)	(0.008)	(0.008)	(0.011)	(0.012)	(0.006)	(0.007)	(0.009)	(0.008)
Observations	2,207	2,207	2,207	2,207	2,207	2,207	2,207	2,207	2,207	2,207	2,207	2,207
Adjusted R^2	0.746	0.751	0.471	0.456	0.672	0.679	0.414	0.396	0.615	0.611	0.589	0.605

Table 2.3: Attention Dynamics

This table shows the OLS regression coefficients for the two time periods lagged attention variables by subgroup. Arellano cluster-robust standard errors are shown in parentheses. The standard errors are clustered at the artist level. ***, **, and * denote significance at the 0.1%, 1%, and 5% level, respectively.

2.6.2 Price Levels of Contemporary Art

Figure 2.4 plots the natural logarithm of the median realized hammer price per square centimeter for deciles of the 12-month-average pre-auction attention scores by type of centrality measure. The measure for the reputation of an artist very strongly positively correlates with realized hammer prices, with a Pearson correlation coefficient ρ higher than 0.9 in either case.

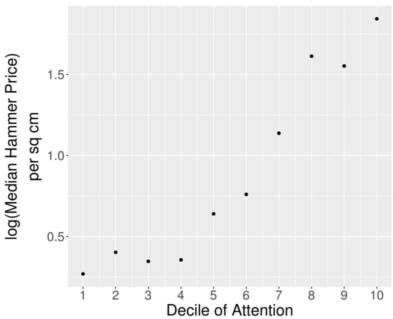

(a) Indegree ($\rho = 0.91$)

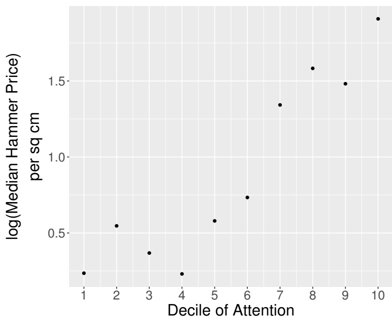

(b) Eigenvector ($\rho = 0.9$)

Figure 2.4: Attention and Prices

This figure shows the log median hammer price per square centimeter as a function of deciles of average status-weighted attention of all art-market subgroups, averaged for the 12 months preceding the auctions. Subfigure 2.4a shows the results when indegree centrality is used for node weighting, and subfigure 2.4b shows the results when eigenvector centrality is used for node weighting.

Table 2.4 shows the results of the hedonic regression of the model described in equation (2.4). The first two columns describe the results when the 12 monthly pre-auction attention scores are not discounted, i.e. $\delta_\tau = 1$, and columns 3-4 describe the results when the monthly attention values are discounted using the data-driven procedure described in Appendix 2.B. Ceteris paribus, a one-standard-deviation increase in status-weighted gallery attention prior to the auction results in an average price increase per square centimeter of about 26% using the averaged coefficients in columns 1-4.[22] The impact of status-weighted gallery attention on the price per square centimeter of an artist is larger than for any other art-market subgroup. This result is consistent with the findings of the attention-dynamics analysis in Table 2.3, which crystallized the art-gallery institution as the key belief generator in the contemporary art market, which initiates and leads the credentialing process of emerging artists. The results of the hedonic regressions show how important it is for the commercial success of an emerging artist to be represented by a high-status gallery.

The subgroup of art advisors exhibits the second largest effect of status-weighted attention on prices. A one-standard-deviation increase in the status-weighted attention of art advisors increases prices by about 17%. The third largest effect of attention on prices is seen with the media subgroup. A one-standard-deviation increase in media attention is expected to increase prices by 11%. The effect concerning the impact of curator credentialing on prices is ambiguous. While the results in Table 2.4 suggest a small positive impact on prices, the robustness tests with differing lagged price and fixed-effect specifications (Table 2.7, 2.8, 2.9) defy a clear interpretation. However, the majority of the results of this

[22] The effect size is determined using the following formula: $100 \cdot (e^{\beta} - 1)$. In addition, all interpretations are made under the ceteris paribus assumption and reflect average effects. Unless otherwise stated, the effect is determined by averaging the coefficients in columns 1-4.

paper support the idea that curator credentialing has no effect on prices.[23]

Moreover, I find no consistent statistically significant effect of museum attention on prices. Similarly, no clear effect of the attention of the outsider group on prices is found. Finally, the effect of status-weighted collector attention on prices per square centimeter is statistically and economically significant. However, the magnitude of the coefficients varies considerably. A one-standard-deviation increase in status-weighted collector attention is expected to increase prices per square centimeter by about 10%. As discussed in Section 2.4, this must be interpreted as the effect of both credentialing and collector demand.[24]

The coefficients for the size of an artwork indicate that the price per square centimeter first decreases with length and width, then increases. This is consistent with the results of other studies that find that, on average, the absolute price of an artwork increases with its size, except

[23] I am grateful to an anonymous reviewer for highlighting the possibility that the finding of curators not influencing prices may be attributable to the fact that many curators in my sample are museum curators, which tend to grant attention to those artists whose prices have already been established. This suggests that there could also be taste-maker curators, which in fact have the potential to impact prices. This discussion indicates that future investigations may benefit from establishing multiple curator categories in order to discern the effects at a more granular level.

[24] In Appendix 2.F the model is re-estimated including artist-fixed effects. While the qualitative nature of the results remains untouched, the effect concerning art advisor credentialing vanishes. An empirical phenomenon attributable to the fact that the effects in this set up are estimated from within artist variation, which is limited due to the short time span of the analysis. The Tables 2.8, 2.9 show the results with alternative lagged price specifications. The qualitative nature of the results remains untouched.

for very large artworks.(Renneboog & Spaenjers, 2013).[25] Moreover, on average paintings by women had a premium of 8.0%. However, this effect is not estimated consistently throughout the robustness tests with alternative lagged price specifications as shown in Appendix 2.G. Paintings sold at Philipps, Christie's, and Sotheby's were on average 55%, 79%, 74% more expensive, respectively, than those sold at other auction houses. On average, paintings by older artists sold for more than those of younger artists. Increasing an artist's age by one year resulted in a price increase of about 2%. The opposite effect was found for the age of artworks. The price of an artwork decreases by about 1% when it becomes one year older. The negative effect of an artwork's age on prices is likely a peculiarity of the market for emerging artists. Concerning, liquidity, I find like Velthuis (2005) a positive effect of liquidity on realized prices. Ceteris paribus, an increase of one painting sold more in the previous month increases prices on average by around 11%. Lastly, I find that past prices affect current prices only marginally. An increase of 1 USD (2017) in the mean realized price per square centimeter in the previous year is expected to raise the current year prices per square centimeter by approximately 2%.

[25] For example, a 5cm X 5cm painting may cost 100 USD, which corresponds to a price per square centimeter of 4 USD. A 10 cm X 10 cm painting may cost 300 USD, which is an absolutely higher price, but the price per square centimeter has dropped to 3 USD.

	log(Hammer Price per Square Centimeter)			
	Undiscounted Attention		Discounted Attention	
	(Eigenvector)	*(Indegree)*	*(Eigenvector)*	*(Indegree)*
A12 Gallery	0.219***	0.237***	0.233***	0.248***
	(0.030)	(0.028)	(0.029)	(0.027)
A12 Art Advisor	0.154***	0.171***	0.143***	0.160***
	(0.032)	(0.030)	(0.029)	(0.028)
A12 Media	0.131***	0.111***	0.093***	0.081***
	(0.025)	(0.024)	(0.023)	(0.023)
A12 Curator	0.045*	0.034	0.066***	0.048**
	(0.025)	(0.023)	(0.025)	(0.024)
A12 Museum	−0.006	−0.026	0.004	−0.013
	(0.024)	(0.024)	(0.023)	(0.023)
A12 Other	−0.117***	−0.034	−0.099***	−0.018
	(0.035)	(0.033)	(0.032)	(0.030)
A12 Collector	0.122***	0.060*	0.122***	0.063**
	(0.033)	(0.031)	(0.031)	(0.029)
Length	−0.642***	−0.645***	−0.642***	−0.646***
	(0.109)	(0.107)	(0.109)	(0.108)
Length SQ	0.244**	0.244**	0.246**	0.247**
	(0.103)	(0.101)	(0.104)	(0.102)
Width	−0.442***	−0.442***	−0.444***	−0.444***
	(0.078)	(0.078)	(0.078)	(0.078)
Width SQ	0.227***	0.226***	0.231***	0.229***
	(0.066)	(0.066)	(0.066)	(0.066)
Female	0.072*	0.088**	0.065*	0.081**
	(0.039)	(0.039)	(0.038)	(0.038)
Phillips	0.438***	0.435***	0.434***	0.433***
	(0.052)	(0.052)	(0.052)	(0.052)
Christie's	0.584***	0.582***	0.582***	0.581***
	(0.054)	(0.054)	(0.054)	(0.054)
Sotheby's	0.555***	0.552***	0.553***	0.550***
	(0.055)	(0.055)	(0.055)	(0.055)
Age Artist	0.020***	0.020***	0.021***	0.020***
	(0.001)	(0.001)	(0.001)	(0.001)
Age Artwork	−0.008***	−0.008***	−0.008***	−0.008***
	(0.002)	(0.002)	(0.002)	(0.002)
Price SQCM T-1 Mean	0.015***	0.015***	0.016***	0.015***
	(0.001)	(0.001)	(0.001)	(0.001)
Liquidity	0.110***	0.107***	0.107***	0.104***
	(0.012)	(0.012)	(0.012)	(0.012)
Constant	−1.219***	−1.254***	−1.163***	−1.199***
	(0.307)	(0.308)	(0.310)	(0.310)
Material FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	5,327	5,327	5,327	5,327
Adjusted R^2	0.555	0.558	0.557	0.559

Table 2.4: Paintings: Art Market Subgroup Attention Evaluation

This table shows the OLS regression results for 4 different hedonic regressions. The first two columns show the results when monthly attention scores are not discounted ($\delta_\tau = 1$)), and columns 3-4 reflect the results when the monthly attention scores are discounted according to the procedure derived in Appendix 2.B. The attention scores are standardized. Arellano cluster-robust standard errors are shown in parentheses. The standard errors are clustered at the artist level. ***, **, and * denote significance at the 0.1%, 1%, and 5% level, respectively.

Overall, special attention is warranted to compare the results concerning the effects of artist reputation on prices to the previous literature. The differences to the seminal studies by Velthuis (2005) and Beckert and Rössel (2013) are of particular relevance. It is important to note that the different studies use different approaches to measure and define artist reputation. Velthuis (2005) discusses the reputational effects from the perspective of institutional recognition.[26] The most striking difference is that he finds many insignificant effects of gallery characteristics on prices and hence concludes that gallery representation does not add economic value, while the results in this study suggests that gallery representation influences prices the most with respect to the reputation-capturing variables. The reason for this result is likely due to the particular choice of variables to measure the reputational effects of galleries on prices in his study. Furthermore, Velthuis (2005) claims to find that museum representation has a positive effect on prices. However, the coefficient of the main variable indicating whether an artist sold artworks to a museum is insignificant. Finally, the coefficients of the variables describing prizes in form of grants defy interpretability. Grants carry a negative sign and the coefficient of larger grants is more negative than that of smaller grants. To sum up, Velthuis (2005) does not find conclusive evidence for the effect of institutional recognition and thus artist reputation on prices. This is in contrast to the empirical results above, which indicate strong positive and statistically significant impact

[26] Velthuis (2005) also regards the positive effect of age on prices as a kind of reputation effect because older artists had more time to spread their name in the community.

of institutional recognition on prices, most notably from the gallery space.

Moreover, the results of Beckert and Rössel (2013) are closer to those of this analysis. They find a statistically significant and positive effect of media awareness and awards on auction prices. However, they do not find a statistically significant reputation effect from gallery representation on prices. The reason for this is most likely the differing manner to quantify gallery representation. Finally, both Velthuis (2005) and Beckert and Rössel (2013) do not determine reputation effects by the other art-market subgroups because of data availability.

2.6.3 Disentangling Credentialing Effects

The impact of artist credentialing varies widely across the different art market subgroups. In descending order of impact, a one-standard-deviation increase in status-weighted gallery, art advisor, or media attention has a positive effect on price per square centimeter of an artwork. No effect of artist credentialing by curators and museums on prices was found. The results can be rationalized with respect to price-attention dynamics and the interest structure of art collectors.

First, a plausible reason that credentialing by an art-market subgroup may have no effect on artwork prices is if that subgroup begins its credentialing only after prices of an artist's artworks have already risen. In other words, the more the intensity of credentialing by a subgroup increases with an artist's realized prices, the more the credentialing of that subgroup can be interpreted as the flagging of successful artists. The hypothesis that museums in particular flag successful artists is supported in Figure 2.5. This plot depicts the annual averages of the ratio between the median status-weighted artist attention in the low and medium price groups and the median status-weighted attention in the high price group. For the low and medium price groups, the aforementioned ratio is lowest for the subgroup of museums. This reveals that the museum subgroup grants relatively more status-weighted attention to commercially successful artists compared to all other subgroups. From this point of view, museums flag commercially successful artists more than the other subgroups. This is consistent with the overall status-weighted attention inequality across all artists, described with a Gini status-weighted attention inequality index per art-market subgroup, as depicted in Figure 2.6. The museum subgroup has the highest Gini coefficient, which reflects that it distributes its status-weighted attention most unevenly across artists in comparison to the other subgroups. The idea that this unequal distribution may be due to a more skewed distribution of museums' centrality values compared to the other subgroups is not supported by the centrality-measure data shown in Figure 2.2.

Art galleries take the opposite position. The average ratios between the median status-weighted attention levels of the low and medium price groups and the median status-weighted attention level in the high price group are the highest for art galleries. Art galleries typically have long-term contractual ties with the artists they represent, and they maximize the prices of their artists (Velthuis, 2005). As such, they are responsible for marketing and promoting all of their artists. Consequently, the differences in status-weighted attention between artists should be mainly due to differences in the statues of the galleries that represent the artists, rather than which of their artists the galleries pay relatively more attention to. This notion is supported by the finding that, on average, 91% of artists received attention from art galleries in a given year, but only 58% of artists received attention from museums.

The other art-market subgroups fall between the price-attention ratio extremes of art galleries and museums. The relative order is consistent with the relative influence of credentialing by the different subgroups on prices found in the hedonic regressions in Table 2.4. The only exception is the price-attention ratio of the media subgroup in the low price group. Furthermore, credentialing by an art market subgroup should also contribute more to establishing an artist's credibility in the eyes of an art collector, the more, ex ante, the art collector is willing to have his or her credibility assessment be influenced by a certain subgroup. Analysis of the interest structure of art collectors presents a possibility to approximate this willingness.

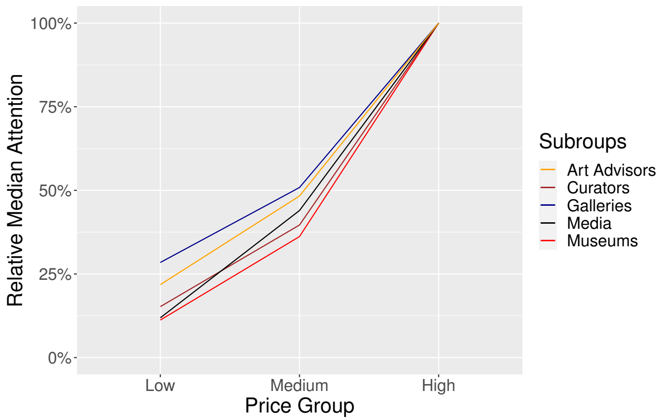

Figure 2.5: Relative Attention and Prices

This figure shows the ratios between the median status-weighted attention level in the low and medium price groups and the median status-weighted attention level in the high price group. For every year between 2015-2021, the artists were separated into the three price groups: high, medium, and low. The 33^{rd} and 66^{th} percentile of the median price per square centimeter of the artists in a given year were used as cut-off values. The plot depicts the average of the yearly ratios.

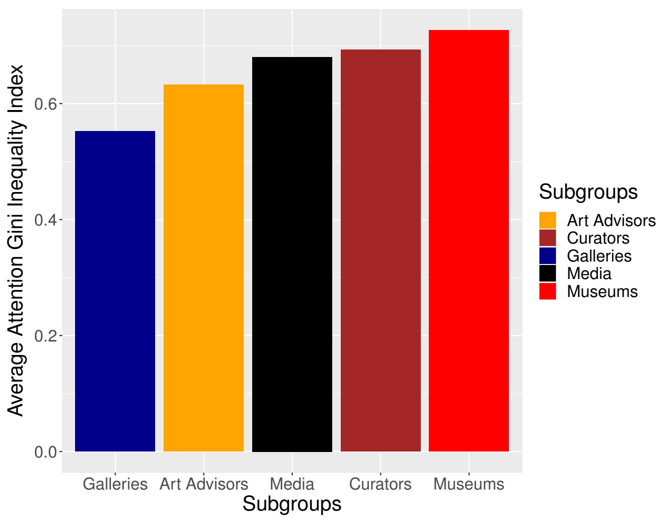

Figure 2.6: Gini Status-Weighted Attention Inequality Index

This figure shows the annual average Gini status-weighted attention inequality index per art market subgroup. The index describes how equally the status-weighted attention is distributed among the artists. A value of 1 describes total inequality.

In order to determine the interest structure of contemporary art collectors, the followings of the accounts representing the art collector subgroups were analyzed.[27] All followings were grouped according to the different subgroups of the art market. For each account in a subgroup, the number of followers was counted, which reflects an indegree centrality statistic of graph theory. For each subgroup, the average indegree of all accounts was measured. The same procedure was repeated for the remaining representatives of the other subgroups of the art market. Figure 2.7 shows that the representatives of the art-collectors subgroup have a different interest-structure pattern than the general art market. On average, art collectors show the most interest in the Instagram accounts of art galleries, followed by those of art advisors, media, curators, and lastly, museums. The representatives of the other subgroups, labelled general art market, reveal a different interest structure. Museums attract the most interest on average, which speaks to their overall importance as an art-market institution. Art advisors receive the least average interest from the representatives of the general art market.

[27] On Instagram an account can follow other accounts. The latter are referred to as followings of the original account.

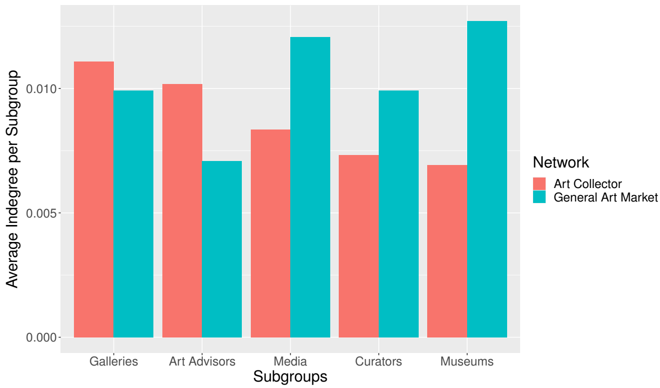

Figure 2.7: Interest Structures

This figure shows the normalized average indegree per subgroup for the art collector and the general art-market network.

In summary, the declining relative influence of art-gallery, art-advisor, and media credentialing processes on artwork prices, and the finding that curator credentialing and museum credentialing are not, on average, price-determining are consistent with the price-attention dynamics results and the interest-structure analysis of the different subgroups.[28] Investigating the interactions of price-attention dynamics and the different interest structures within the hedonic framework – although beyond the scope of this paper – presents an interesting avenue for future research.

[28] I am grateful to an anonymous reviewer for highlighting the possibility that the finding of curators not influencing prices may be attributable to the fact that many curators in my sample are museum curators, which tend to grant attention to those artists whose prices have already been established. This suggests that there could also be taste-maker curators, which in fact have the potential to impact prices. This discussion indicates that future investigations may benefit from establishing multiple curator categories in order to discern the effects at a more granular level.

2.7 Conclusion

This paper developed and applied a methodology that dynamically quantifies reputation of artists. The novel Instagram data set offers the possibility to identify the sources of an artist's reputation at different subgroup levels of the art market. In addition, the methodology permits one to analyze the relative influence of different art-market intermediaries on the economic value creation of products for which there are no clear objective quality standards. In an application for the market of emerging contemporary artists, it was shown that art galleries act as gatekeepers. They initiate and lead the credentialing process of emerging artists; as such, art galleries comprise the key belief-generating institution in this market. Furthermore, it was shown that an artist's credentialing by art galleries drives the artist's prices the most. This reveals the fundamental importance for the commercial success of emerging artists to be represented by a high-status art gallery. Artist credentialing by art advisors and the media was found to also significantly impact artist prices. No effect of artist credentialing by museums and curators on prices was found. The relative effects of credentialing by different art-market subgroups on prices was rationalized with respect to price-attention dynamics of the various art-market subgroups and the interest structure of art collectors. The price-attention dynamics analysis revealed that those subgroups for which no effect of the credentialing on prices could be established have a tendency to grant already commercially successful artists relatively more status-weighed attention compared to the other art-market subgroups.

In other words, those subgroups tend to flag already commercially successful artists. Moreover, the subgroups in which art collectors show the most interest are more likely to influence art collectors' willingness to pay through their artist credentialing. Analysis of the interest structure of the representatives of the art-collector subgroup revealed that they are primarily interested in the undertakings of art galleries, followed by undertakings of art advisors, media, curators, and museums. The art collector-interest structure analysis further rationalized the empirical results regarding the relative influence of different art-market subgroups on artist prices. The introduced methodology for dynamically quantifying an artist's reputation adds an important set of hedonic variables to hedonic art-pricing models. This study focused on emerging artists. Future research opportunities include applying the methodology to other, more or less established, artist groups as well. Furthermore, the methodology can also be readily adapted to better understand other markets for goods which lack objective quality standards.

References

Accominotti, F. (2022). *The price of hierarchy: Consecration and the formation of economic value in the market for modern art.* (Working Paper)

Anand, N., & Jones, B. C. (2008). Tournament rituals, category dynamics, and field configuration: The case of the booker prize. *Journal of Management Studies*, *45*, 1036-1060.

Antal, A. B., Hutter, M., & Stark, D. (2015). *Moments of valuation: Exploring sites of dissonance.* Oxford: Oxford University Press.

Aubry, M., Kräussl, R., Manso, G., & Spaenjers, C. (2023). Biased auctioneers. *Journal of Finance*, 1-43.

Baumol, W. J. (1986). Unnatural value: Or art investment as floating crap game. *American Economic Review*, *76*(2), 10-14.

Becker, H. S. (1982). *Art worlds.* Berkeley and Los Angeles: University of California Press.

Beckert, J. (2020). Markets from meaning: quality uncertainty and the intersubjective construction of value. *Cambridge Journal of Economics*, *44*(1), 285-301.

Beckert, J., & Aspers, P. (2011). *The worth of goods: Valuation and pricing in the economy.* New York: Oxford University Press.

Beckert, J., & Rössel, J. (2004). Kunst und Preise - Reputation als Mechanismus der Reduktion von Ungewissheit am Kunstmarkt. *Kölner Zeitschrift für Soziologie und Sozialpsychologie*, *56*(1), 32-50.

Beckert, J., & Rössel, J. (2013). The price of art: Uncertainty and reputation in the art field. *European Societies*, *15*(2), 178-195.

Beljean, S., Chong, P., & Lamont, M. (2015). A post-Bourdieusian sociology of valuation and evaluation for the field of cultural production. In L. Hanquinet & M. Savage (Eds.), *Routledge International Handbook of the Sociology of Art and Culture* (p. 38-48). London: Routledge.

Boltanski, L., & Thévenot, L. (2006). *On justification: Economies of worth.* Princeton, NJ: Princeton University Press.

Bonus, H., & Ronte, D. (1997). Credibility and economic value in the visual arts. *Journal of Cultural Economics*, *21*(1), 103-118.

Bourdieu, P. (1993). *The field of cultural production.* New York: Columbia University Press.

Bourdieu, P. (1996). *The rules of art: Genesis and structure of the literary field.* Cambridge, MA: Polity Press.

Bystryn, M. (1978). Art galleries as gatekeepers: The case of the abstract expressionists. *Social Research*, *45*(2), 390-408.

Campos, N. F., & Barbosa, R. L. (2009). Paintings and numbers: An econometric investigation of sales rates, prices, and returns in latin american art auctions. *Oxford Economic Papers*, *61*(1), 28-51.

Chanel, O., & Gérard-Varet, L.-A. (1996). The relevance of hedonic price indices. *Journal of Cultural Economics*, *20*(1), 1-24.

Coslor, E., Crawford, B., & Leyshon, A. (2019). Collectors, investors and speculators: Gatekeeper use of audience categories in the art market. *Organization Studies*, *41*(7), 945–967.

de la Barre, M., Docclo, S., & Ginsburgh, V. (1994). Returns of impressionist, modern and contemporary European paintings 1962-1991. *Annalyses d'Economie et de Statistique*, *35*(1), 143-181.

DiMaggio, P., & Powell, W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*.

English, J. F. (2005). *The economy of prestige: Prizes, awards, and the circulation of cultural value.* Cambridge, MA: Harvard University Press.

Etro, F., & Pagani, L. (2013). The market for paintings in the Venetian Republic from Renaissance to Rococò. *Journal of Cultural Economics*, *37*(4), 391-415.

Fligstein, N., & McAdam, D. (2012). *A Theory of Fields.* Oxford, Oxford University Press.

Foster, P., Borgatti, S. P., & Jones, C. (2011). Gatekeeper search and selection strategies: Relational and network governance in a cultural market. *Poetics*, *39*(4), 247–265.

Fraiberger, S., Sinatra, R., Resch, M., Riedl, C., & Barabasi, A. (2018). Quantifying reputation and success in art. *Science*, *362*(6416), 825–829.

Frey, B. S., & Pommerehne, W. W. (1989). *Muses and markets: Explorations in the economics of the arts.* Basil Blackwell, Oxford.

Godart, F., Hsu, G., & Negro, G. (2023). Gatekeeping and the use of contested practices in creative industries: The case of fur in fashion. *Organization Science*, *34*(2), 637–656.

Gould, R. V. (2002). The origins of status hierarchies: A formal theory and empirical test. *American Journal of Sociology*, *107*(5), 1143–1178.

Grampp, W. D. (1989). *Pricing the priceless: Art, artists, and economics.* Basic Books, New York.

Jackson, M. O. (2011). *Social and economic networks.* Princeton University Press.

Karpik, L. (2010). *Valuing the unique: The economics of singularities.* Princeton University Press.

Kawashima, N. (1999). Distribution of the arts: British arts centres as 'gatekeepers' in intersecting cultural production systems. *Poetics*, *26*(4), 263–283.

Khaire, M., & Wadhwani, R. D. (2010). Changing landscapes: The construction of meaning and value in a new market category—modern Indian art. *Academy of Management Journal*, *53*(6), 1281–1304.

Lamont, M. (1987). How to become a dominant French philosopher: The case of Jacques Derrida. *American Journal of Sociology*, *93*(3), 584-622.

Ma, M. X., Noussair, C. N., & Renneboog, L. (2022). Colors, emotions, and the auction value of paintings. *European Economic Review*, *142*, 1-26.

Marinelli, N., & Palombo, G. (2011). A model for pricing Italian contemporary art paintings at auction. *The Quarterly Review of Economics and Finance*, *51*(2), 212-224.

Menger, P.-M. (2014). *The economics of creativity: Art and achievement* *under uncertainty.* Cambridge, MA: Harvard University Press.

Mulkay, M., & Chaplin, E. (1982). Aesthetics and the artistic career: A study of anomie in fine-art painting. *The Sociological Quarterly*, *23*(1), 117-138.

Podolny, J. M. (1993). A status-based model of market competition. *American Journal of Sociology*, *89*(4), 829-872.

Rengers, M., & Velthuis, O. (2002). Determinants of prices for contemporary art in Dutch galleries. *Journal of Cultural Economics*, *26*, 1-28.

Renneboog, L., & Spaenjers, C. (2011). The iconic boom in modern Russian art. *The Journal of Alternative Investments*, *13*(3), 67-80.

Renneboog, L., & Spaenjers, C. (2013). Buying beauty: On prices and returns in the art market. *Management Science*, *59*(1), 36-53.

Salganik, M. J., Dodd, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, *311*, 854-856.

Smith, B. H. (1988). *Contingencies of value.* Cambridge, MA: Harvard University Press.

Smits, R. (2016). Gatekeeping and networking arrangements: Dutch distributors in the film distribution business. *Poetics*, *58*, 29–42.

Sproule, R., & Valsan, C. (2006). Hedonic models and pre-auction estimates: Abstract art revisted. *Economics Bulletin*, *26*(5), 1-10.

Stark, D. (2009). *The sense of dissonance: Accounts of worth in economic life.* Princeton; NJ: Princeton University Press.

Taylor, D., & Coleman, L. (2011). Price determinants of Aboriginal art, and its role as an alternative asset class. *Journal of Banking & Finance*, *35*(6), 1519-1529.

Ullrich, W. (2017, July). Zwischen Deko und Diskurs [Between deco and discourse]. *Perlentaucher*. (`https://www.perlentaucher.de/essay/wolfgang-ullrich-ueber-kuratoren-und-kunstmarktkunst.html` Retrieved May, 2022)

Vasan, K., Janosov, M., & Barabási, A.-L. (2022). Quantifying NFT-driven networks in crypto art. *Scientific Reports*, *12*(1).

Velthuis, O. (2005). *Talking prices: Symbolic meaning of prices on the market for contemporary art.* Princeton, NJ: Princeton University Press.

Appendix

2.A Top 100 Gallery and Museum Instagram Usernames

- 1. artbasel
- 2. artforum
- 3. tate
- 4. themuseumofmodernart
- 5. frieze_magazine
- 6. friezeartfair
- 7. artnet
- 8. guggenheim
- 9. davidzwirner
- 10. artsy
- 11. hauserwirth
- 12. hansulrichobrist
- 13. whitneymuseum
- 14. gagosian
- 15. momaps1
- 16. newmuseum
- 17. artnews
- 18. centrepompidou
- 19. metmuseum
- 20. mariangoodmangallery
- 21. whitecube
- 22. pacegallery
- 23. jerrysaltz
- 24. artreview_magazine
- 25. palaisdetokyo
- 26. moussemagazine
- 27. theartnewspaper.official
- 28. collecteurs
- 29. diaartfoundation
- 30. moca
- 31. serpentineuk
- 32. lisson_gallery
- 33. flashartmagazine
- 34. fiacparis

- 35. brooklynmuseum
- 36. lacma
- 37. artobserved
- 38. blouin_artinfo
- 39. klausbiesenbach
- 40. fondationbeyeler
- 41. sothebys
- 42. labiennale
- 43. thearmoryshow
- 44. thaddaeusropac
- 45. fondazioneprada
- 46. artinamerica
- 47. koeniggalerie
- 48. artspace
- 49. simondepury
- 50. christiesinc
- 51. thecollectorslist
- 52. galerieperrotin
- 53. whitechapelgallery
- 54. hammer_museum
- 55. victoriamirogallery
- 56. sfmoma
- 57. paulacoopergallery
- 58. sadiecoleshq
- 59. hyperallergic
- 60. alminerech
- 61. c.a.daily
- 62. guggenheim_venice
- 63. stuartcomer
- 64. levygorvy
- 65. saatchi_gallery
- 66. simonleegallery
- 67. spruethmagers
- 68. independent_hq
- 69. galleriesnow
- 70. xavierhufkens
- 71. petzelgallery
- 72. stedelijkmuseum
- 73. museelouvre
- 74. drawingcenter
- 75. anitazart
- 76. artbrussels
- 77. e_flux
- 78. britishmuseum
- 79. nespector
- 80. art.viewer
- 81. davidkordanskygallery

- 82. flagartfoundation
- 83. massimodecarlogallery
- 84. fondationlv
- 85. 303gallery
- 86. ceciliaalemani
- 87. museoguggenheim
- 88. vamuseum
- 89. icalondon
- 90. phillipsauction
- 91. walkerartcenter
- 92. the_adaa
- 93. gavinbrownsenterprise
- 94. rubellmuseum
- 95. bourgeois0490
- 96. newartdealers
- 97. kunstsammler
- 98. artissimafair
- 99. galeriemaxhetzler
- 100. larrys_list

2.B Data-Driven Discount Factors

In order to account for the possibility that attention closer to the auction date influences hammer prices more, I estimate the following model using OLS:

$$ln(p_{j,t}) = \alpha + \sum_{\tau=1}^{12} \delta_{\tau}\, \mathrm{A}_{i,t-\tau} + \sum_{m=1}^{M} \beta_m \mathrm{Artist}_{m,j,t} + \sum_{n=1}^{N} \beta_n \mathrm{Transaction}_{n,j,t} + \sum_{p=1}^{P} \beta_p \mathrm{Physical}_{p,j,t} + \epsilon_{j,t}, \tag{2.5}$$

where $A_{i,t-\tau}$ describes the status-weighted attention score of artist i of all subgroups, τ months before the auction. Since the relative influence of different lags of status-weighted attention on prices only marginally depends on the choice of centrality measure used to approximate the node weights, the averages of the coefficients per lag τ are used as the data-driven discount factors δ_{τ} in model (2.4).

The regression output concerning the coefficients δ_{τ} is depicted in Table 2.5. The regression coefficients δ_{τ} are used to discount the respective monthly attention scores in a data-driven manner:

	log(Hammer Price per Square Centimeter)	
	Eigenvector	Indegree
L1 Attention	0.134***	0.127***
	(0.023)	(0.023)
L2 Attention	0.073***	0.081***
	(0.026)	(0.026)
L3 Attention	0.005	0.005
	(0.021)	(0.021)
L4 Attention	0.078***	0.081***
	(0.022)	(0.022)
L5 Attention	0.038	0.046*
	(0.028)	(0.027)
L6 Attention	0.057***	0.056***
	(0.021)	(0.021)
L7 Attention	0.036	0.038
	(0.026)	(0.027)
L8 Attention	0.090***	0.084***
	(0.025)	(0.024)
L9 Attention	0.093***	0.100***
	(0.025)	(0.024)
L10 Attention	0.038	0.041
	(0.029)	(0.028)
L11 Attention	0.012	0.009
	(0.027)	(0.027)
L12 Attention	0.055**	0.048*
	(0.026)	(0.026)
Material FE	Yes	Yes
Year FE	Yes	Yes
Month FE	Yes	Yes
Observations	5,269	5,269
Adjusted R^2	0.543	0.547

Table 2.5: Data-Driven Identification of Discount Factors

This table shows the OLS regression results of model (2.5). The attention scores are standardized. Arellano cluster-robust standard errors are shown in parentheses. The standard errors are clustered at the artist level. ***, **, and * denote significance at the 0.1%, 1%, and 5% level, respectively.

2.C Relative Importance Subgroup Representatives

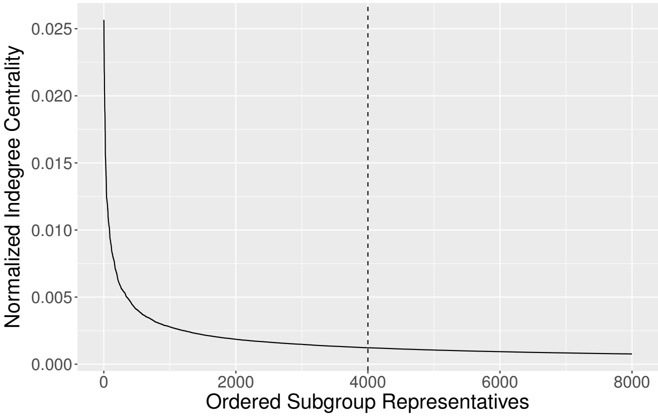

Figure 2.8: Relative Importance of Subgroup Representatives

This figure shows the exponential decay of the indegree centrality statistic of the top 8,000 art market subgroup representatives.

2.D Curator Attention and Prices

Ullrich (2017) advances the idea that commercially successful artists might be denied the ability to empathize with minorities or underprivileged milieus, and thus to make art that meets the sociopolitical demands of curators. A likely consequence would be that commercially successful artists receive less status-weighted curator attention compared to the less commercially successful artists. Figure 2.9 contradicts this notion.

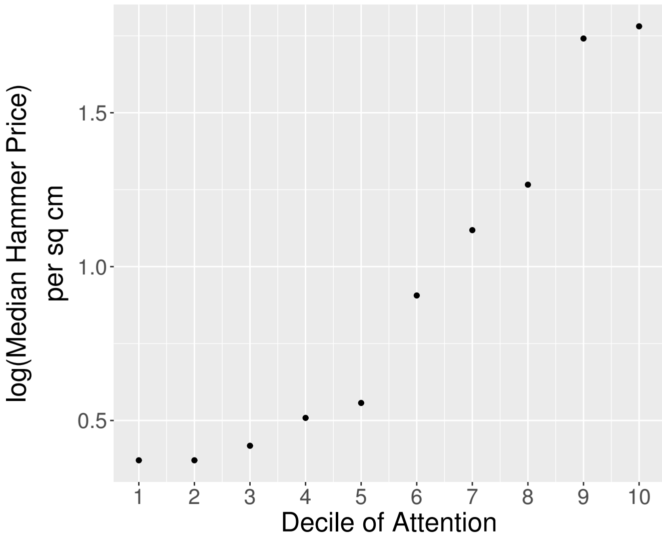

Figure 2.9: Prices and Curator Attention

This figure shows the log median hammer price per square centimeter as a function of the deciles of the average 12-month pre-auction status-weighted curator attention using indegree centrality for node weighting. The Pearson correlation coefficient is $\rho = 0.91$.

2.E Random Subset Robustness Test

	log(Hammer Price per Square Centimeter)			
	Undiscounted Attention		Discounted Attention	
	(Eigenvector)	*(Indegree)*	*(Eigenvector)*	*(Indegree)*
A12 Gallery	0.185^{***}	0.216^{***}	0.211^{***}	0.236^{***}
	(0.032)	(0.031)	(0.032)	(0.030)
A12 Art Advisor	0.202^{***}	0.199^{***}	0.193^{***}	0.192^{***}
	(0.036)	(0.034)	(0.034)	(0.032)
A12 Media	0.104^{***}	0.100^{***}	0.062^{**}	0.067^{***}
	(0.028)	(0.027)	(0.027)	(0.025)
A12 Curator	0.019	0.004	0.031	0.010
	(0.026)	(0.024)	(0.027)	(0.026)
A12 Museum	0.024	−0.002	0.038	0.013
	(0.026)	(0.026)	(0.025)	(0.025)
A12 Other	-0.107^{***}	−0.020	-0.097^{***}	−0.014
	(0.038)	(0.035)	(0.034)	(0.032)
A12 Collector	0.113^{***}	0.048	0.115^{***}	0.054^{*}
	(0.035)	(0.033)	(0.033)	(0.031)
Length	-0.693^{***}	-0.688^{***}	-0.692^{***}	-0.688^{***}
	(0.116)	(0.114)	(0.116)	(0.115)
Length SQ	0.266^{**}	0.263^{**}	0.268^{**}	0.265^{**}
	(0.111)	(0.109)	(0.111)	(0.109)
Width	-0.400^{***}	-0.404^{***}	-0.404^{***}	-0.407^{***}
	(0.082)	(0.081)	(0.082)	(0.081)
Width SQ	0.187^{***}	0.186^{***}	0.193^{***}	0.191^{***}
	(0.068)	(0.067)	(0.068)	(0.067)
Female	0.023	0.031	0.021	0.030
	(0.040)	(0.040)	(0.040)	(0.040)
Phillips	0.463^{***}	0.458^{***}	0.457^{***}	0.455^{***}
	(0.056)	(0.056)	(0.056)	(0.056)
Christie's	0.596^{***}	0.596^{***}	0.594^{***}	0.593^{***}
	(0.059)	(0.059)	(0.059)	(0.059)
Sotheby's	0.581^{***}	0.579^{***}	0.581^{***}	0.578^{***}
	(0.060)	(0.060)	(0.060)	(0.060)
Age Artist	0.022^{***}	0.021^{***}	0.022^{***}	0.021^{***}
	(0.002)	(0.001)	(0.001)	(0.001)
Age Artwork	-0.009^{***}	-0.009^{***}	-0.009^{***}	-0.009^{***}
	(0.002)	(0.002)	(0.002)	(0.002)
Price SQCM T-1 Mean	0.015^{***}	0.015^{***}	0.015^{***}	0.015^{***}
	(0.001)	(0.001)	(0.001)	(0.001)
Liquidity	0.122^{***}	0.118^{***}	0.120^{***}	0.116^{***}
	(0.013)	(0.013)	(0.013)	(0.013)
Constant	-1.341^{***}	-1.364^{***}	-1.302^{***}	-1.319^{***}
	(0.313)	(0.313)	(0.314)	(0.314)
Material FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	4,639	4,639	4,639	4,639
Adjusted R^2	0.561	0.563	0.563	0.565

Table 2.6: Paintings: Art Market Subgroup Attention Evaluation

This table shows the OLS regression results for 4 different hedonic regressions. The first two columns show the results when monthly attention scores are not discounted ($\delta_\tau = 1$)), and columns 3-4 reflect the results when the monthly attention scores are discounted according to the procedure derived in Appendix 2.B. The attention scores are standardized. Arellano cluster-robust standard errors are shown in parentheses. The standard errors are clustered at the artist level. ***, **, and * denote significance at the 0.1%, 1%, and 5% level, respectively.

2.F Artist Fixed Effects Price Regression

	log(Hammer Price per Square Centimeter)			
	Undiscounted Attention		Discounted Attention	
	(Eigenvector)	*(Indegree)*	*(Eigenvector)*	*(Indegree)*
A12 Gallery	0.086^{**}	0.092^{**}	0.114^{***}	0.116^{***}
	(0.038)	(0.036)	(0.035)	(0.034)
A12 Art Advisor	0.003	−0.016	−0.004	−0.019
	(0.033)	(0.031)	(0.029)	(0.027)
A12 Media	0.076^{***}	0.073^{***}	0.043^{**}	0.045^{**}
	(0.024)	(0.023)	(0.021)	(0.020)
A12 Curator	−0.003	−0.003	0.022	0.014
	(0.033)	(0.032)	(0.028)	(0.028)
A12 Museum	-0.048^{*}	-0.046^{*}	−0.016	−0.017
	(0.025)	(0.024)	(0.024)	(0.023)
A12 Other	−0.054	−0.014	−0.044	−0.007
	(0.040)	(0.037)	(0.034)	(0.031)
A12 Collector	0.161^{***}	0.157^{***}	0.127^{***}	0.125^{***}
	(0.039)	(0.038)	(0.036)	(0.035)
Length	-0.323^{***}	-0.323^{***}	-0.322^{***}	-0.323^{***}
	(0.075)	(0.076)	(0.075)	(0.076)
Length SQ	0.075	0.076	0.075	0.076
	(0.067)	(0.067)	(0.066)	(0.067)
Width	-0.551^{***}	-0.553^{***}	-0.553^{***}	-0.554^{***}
	(0.062)	(0.062)	(0.062)	(0.062)
Width SQ	0.259^{***}	0.261^{***}	0.261^{***}	0.262^{***}
	(0.051)	(0.051)	(0.051)	(0.051)
Phillips	0.348^{***}	0.345^{***}	0.350^{***}	0.349^{***}
	(0.045)	(0.045)	(0.045)	(0.045)
Christie's	0.349^{***}	0.347^{***}	0.350^{***}	0.349^{***}
	(0.047)	(0.047)	(0.048)	(0.048)
Sotheby's	0.333^{***}	0.332^{***}	0.335^{***}	0.335^{***}
	(0.047)	(0.047)	(0.047)	(0.047)
Age Artist	-0.061^{***}	-0.061^{***}	-0.061^{***}	-0.060^{***}
	(0.007)	(0.007)	(0.007)	(0.007)
Age Artwork	-0.007^{***}	-0.007^{***}	-0.007^{***}	-0.007^{***}
	(0.002)	(0.002)	(0.002)	(0.002)
Price SQCM T-1 Mean	0.0002	0.0002	0.0004	0.0004
	(0.001)	(0.001)	(0.001)	(0.001)
Liquidity	0.013	0.012	0.011	0.011
	(0.010)	(0.010)	(0.010)	(0.010)
Constant	6.981^{***}	6.952^{***}	6.990^{***}	6.943^{***}
	(0.566)	(0.567)	(0.563)	(0.565)
Artist FE	Yes	Yes	Yes	Yes
Material FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	5,327	5,327	5,327	5,327
Adjusted R^2	0.770	0.771	0.770	0.770

Table 2.7: Paintings: Art Market Subgroup Attention Evaluation

This table shows the OLS regression results for 4 different hedonic regressions. The first two columns show the results when monthly attention scores are not discounted ($\delta_\tau = 1$)), and columns 3-4 reflect the results when the monthly attention scores are discounted according to the procedure derived in Appendix 2.B. The attention scores are standardized. Arellano cluster-robust standard errors are shown in parentheses. The standard errors are clustered at the artist level. ***, **, and * denote significance at the 0.1%, 1%, and 5% level, respectively.

2.G Alternative Lagged Price Specifications

	log(Hammer Price per Square Centimeter)			
	Undiscounted Attention		Discounted Attention	
	(Eigenvector)	*(Indegree)*	*(Eigenvector)*	*(Indegree)*
A12 Gallery	0.237***	0.256***	0.251***	0.267***
	(0.030)	(0.028)	(0.029)	(0.028)
A12 Art Advisor	0.177***	0.195***	0.163***	0.182***
	(0.032)	(0.030)	(0.030)	(0.028)
A12 Media	0.139***	0.116***	0.101***	0.088***
	(0.026)	(0.025)	(0.024)	(0.023)
A12 Curator	0.038	0.028	0.059**	0.041*
	(0.025)	(0.024)	(0.026)	(0.024)
A12 Museum	−0.020	−0.040	−0.010	−0.028
	(0.024)	(0.024)	(0.023)	(0.023)
A12 Other	−0.124***	−0.037	−0.102***	−0.019
	(0.036)	(0.034)	(0.032)	(0.030)
A12 Collector	0.109***	0.042	0.108***	0.045
	(0.034)	(0.032)	(0.032)	(0.030)
Length	−0.653***	−0.656***	−0.654***	−0.657***
	(0.111)	(0.109)	(0.112)	(0.110)
Length SQ	0.251**	0.250**	0.253**	0.253**
	(0.105)	(0.103)	(0.106)	(0.105)
Width	−0.448***	−0.447***	−0.450***	−0.449***
	(0.079)	(0.078)	(0.079)	(0.079)
Width SQ	0.230***	0.227***	0.234***	0.232***
	(0.066)	(0.066)	(0.067)	(0.066)
Female	0.053	0.069*	0.045	0.061
	(0.039)	(0.038)	(0.038)	(0.038)
Phillips	0.438***	0.434***	0.433***	0.431***
	(0.052)	(0.052)	(0.052)	(0.052)
Christie's	0.584***	0.581***	0.582***	0.579***
	(0.054)	(0.054)	(0.054)	(0.053)
Sotheby's	0.558***	0.554***	0.556***	0.552***
	(0.055)	(0.055)	(0.055)	(0.055)
Age Artist	0.020***	0.020***	0.020***	0.020***
	(0.001)	(0.001)	(0.001)	(0.001)
Age Artwork	−0.007***	−0.008***	−0.007***	−0.008***
	(0.002)	(0.002)	(0.002)	(0.002)
Price SQCM T-1 Median	0.017***	0.017***	0.017***	0.017***
	(0.001)	(0.001)	(0.001)	(0.001)
Liquidity	0.115***	0.112***	0.112***	0.108***
	(0.012)	(0.012)	(0.012)	(0.012)
Constant	−1.166***	−1.201***	−1.105***	−1.141***
	(0.308)	(0.308)	(0.311)	(0.311)
Material FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	5,327	5,327	5,327	5,327
Adjusted R^2	0.548	0.551	0.549	0.552

Table 2.8: Paintings: Art Market Subgroup Attention Evaluation

This table shows the OLS regression results for 4 different hedonic regressions. The first two columns show the results when monthly attention scores are not discounted ($\delta_\tau = 1$)), and columns 3-4 reflect the results when the monthly attention scores are discounted according to the procedure derived in Appendix 2.B. The attention scores are standardized. Arellano cluster-robust standard errors are shown in parentheses. The standard errors are clustered at the artist level. ***, **, and * denote significance at the 0.1%, 1%, and 5% level, respectively.

	log(Hammer Price per Square Centimeter)			
	Undiscounted Attention		Discounted Attention	
	(Eigenvector)	*(Indegree)*	*(Eigenvector)*	*(Indegree)*
A12 Gallery	0.258***	0.271***	0.270***	0.282***
	(0.031)	(0.029)	(0.030)	(0.028)
A12 Art Advisor	0.193***	0.208***	0.171***	0.188***
	(0.033)	(0.031)	(0.031)	(0.029)
A12 Media	0.081***	0.065***	0.053**	0.043*
	(0.026)	(0.025)	(0.024)	(0.023)
A12 Curator	0.070***	0.063***	0.086***	0.072***
	(0.026)	(0.024)	(0.026)	(0.025)
A12 Museum	−0.046*	−0.068***	−0.029	−0.049**
	(0.026)	(0.026)	(0.025)	(0.024)
A12 Other	−0.125***	−0.035	−0.102***	−0.017
	(0.037)	(0.035)	(0.034)	(0.031)
A12 Collector	0.152***	0.084**	0.149***	0.084***
	(0.034)	(0.033)	(0.032)	(0.030)
Length	−0.665***	−0.668***	−0.663***	−0.667***
	(0.110)	(0.108)	(0.111)	(0.110)
Length SQ	0.254**	0.254**	0.255**	0.256**
	(0.104)	(0.102)	(0.105)	(0.103)
Width	−0.487***	−0.485***	−0.492***	−0.490***
	(0.081)	(0.081)	(0.081)	(0.081)
Width SQ	0.254***	0.251***	0.260***	0.256***
	(0.067)	(0.067)	(0.068)	(0.068)
Female	0.055	0.073*	0.046	0.064
	(0.040)	(0.040)	(0.040)	(0.040)
Phillips	0.470***	0.466***	0.467***	0.465***
	(0.055)	(0.055)	(0.054)	(0.054)
Christie's	0.635***	0.632***	0.635***	0.633***
	(0.056)	(0.056)	(0.056)	(0.056)
Sotheby's	0.639***	0.635***	0.639***	0.635***
	(0.058)	(0.058)	(0.058)	(0.058)
Age Artist	0.022***	0.021***	0.022***	0.021***
	(0.001)	(0.001)	(0.001)	(0.001)
Age Artwork	−0.007***	−0.007***	−0.007***	−0.007***
	(0.002)	(0.002)	(0.002)	(0.002)
Price SQCM T-1 Max	0.003***	0.003***	0.003***	0.003***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Liquidity	0.103***	0.100***	0.100***	0.097***
	(0.013)	(0.013)	(0.013)	(0.013)
Constant	−1.248***	−1.276***	−1.196***	−1.226***
	(0.317)	(0.319)	(0.320)	(0.321)
Material FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	5,327	5,327	5,327	5,327
Adjusted R^2	0.516	0.519	0.518	0.520

Table 2.9: Paintings: Art Market Subgroup Attention Evaluation

This table shows the OLS regression results for 4 different hedonic regressions. The first two columns show the results when monthly attention scores are not discounted ($\delta_\tau = 1$)), and columns 3-4 reflect the results when the monthly attention scores are discounted according to the procedure derived in Appendix 2.B. The attention scores are standardized. Arellano cluster-robust standard errors are shown in parentheses. The standard errors are clustered at the artist level. ***, **, and * denote significance at the 0.1%, 1%, and 5% level, respectively.

2.H Additional Instagram Descriptive Statistics

Subgroup	***Median Followers***	***Median Followings***	***Median Conributions***
Galleries	10,726	1,164	836
Curators	5,028	1,324	892
Museums	36,220	622	1,355
Collectors	6,827	1,020	814
Media	30,671	1,134	1,800
Art Advisors	4,746	1,618	1,229
Others	16,243	1,003	907

Table 2.10: Descriptive Statistics by Subgroup

This table shows the different subgroups of art-market players and selected descriptive statistics of their Instagram profiles. Median followers describes the median of how many Instagram accounts follow each subgroup. Median followings shows the median of how many Instagram accounts the subgroups follow themselves. Median contributions describes the median number of contributions (posts) the subgroups have on their own profile.

Chapter 3

How to optimally select contemporary artists from an investment perspective

Abstract

This paper presents an application of quantitative artist reputation measures. These measures reflect how much status-weighted attention different artists receive from a large set of representatives of the field of art. It is shown that the past dynamics of collector artist attention is predictive of future returns at the artist level. The trading application reveals that selecting artists based on their past collector attention dynamics yields returns that are significantly higher in economic and statistical terms compared to a random selection of artists. The proposed methodology also allows mitigation of significant monetary losses when investing in contemporary art by avoiding selection of those artists with relatively bad recent track records of past collector attention compared to other artists.

Keywords: contemporary art; investment; reputation

3.1 Introduction

In cultural markets, such as those for one for wine, fashion or art, objective quality standards are difficult to determine. Market intermediaries shape the perceived quality of cultural products (Beckert, 2020; Bourdieu, 1993). For the market for contemporary art, recent research has revealed that artist credentialing by various art-market subgroups influences art prices (Serwart, 2023). Artist credentialing refers to the bestowal of attention on the artist by reputable actors in the field that generates his or her reputation (Beckert & Rössel, 2013). Serwart (2023) also shows that the impact of artist credentialing on prices varies by art-market subgroup. However, an analysis concerning the returns on art at the artist level, which takes into account credentialing processes is still missing.

With respect to returns on art, the financial art literature has analyzed the risk-return relationship of art — Korteweg et al. (2016); Lovo & Spaenjers (2018); Mei & Moses (2002); Renneboog & Spaenjers (2013) —, and its financial and macro-economic drivers such as equity market evolution and income inequality; see Goetzmann et al. (2011). Other studies have investigated the role of human biases in the art market, such as anchoring in past prices or auction estimates, which affect returns; see for example Beggs and Kathryn (2009). However, the impact of the dynamics of an artist's reputation on his or her returns has largely been neglected. The closest to being an exception is the paper by Pénasse et al. (2014), who use semi-annual sentiment survey data for the short-term confidence in 70 artists on the positive, neutral and negative scale to estimate the effect of confidence sentiment on the return of art. The lack of research on the effect of the dynamics of artist reputation on the return of of art at the artist level stems from the fact that only very recently has a dynamic and quantitative reputation measure been developed (Serwart, 2023).

This paper proceeds as follows: Section 3.2 describes the model of the artist credentialing process. Section 3.3 derives the construction of the quantitative artist reputation measure and specifies the different econometric models. Section 3.4 presents the data. Section 3.5 contains the empirical results of the models and a trading application. Section 3.6 concludes.

3.2 Credentialing Process Model

The model of the credentialing process, which generates an artist's reputation, is shown in Figure 3.2. The model is taken from Serwart (2023). The field of art comprises several subgroups. Each subgroup is made up of different individual actors. An artist's reputation is gained through the status-weighted attention an artist receives from other insiders in the art industry. There are different ways that an artist can gain reputation in the art world. These can include having an exhibition at a reputable art gallery, being featured in an art magazine and being recommended by an art advisor to a potential buyer, among others. The status-weighted attention an artist receives can be analyzed and subdivided based on the specific subgroups within the art field that are giving attention. The reputation is made up of the status-weighted attention they receive from all the different subgroups within the art industry. An artist's reputation is based on the expectation that he or she will create art that will be well-regarded by influential figures in the industry with a certain level of status. The same concept applies to reputations at the subgroup level.[1] Collectors observe the status-weighted attention an artist receives and use it to form their own personal evaluation of the artist's credibility. A collector transforms the status-weighted attention an artist receives from a subgroup based on his or her own subgroup-specific impact factor. The impact factor indicates the degree to which a collector's assessment of an artist's credibility is influenced by status-weighted attention the artist receives from a specific subgroup within the art market. A collector's evaluation of an artist's credibility assessed in this way determines the collector's belief in the overall artistic quality of that artist, which drives the collector's reputation-based willingness to pay for said artist. Factors specific to a particular piece of art, such as its size, material, and age, ultimately influence an art collector's willingness to pay for it.

[1]The ideas and concepts discussed thus far in this model are based on Beckert (2020); Beckert and Rössel (2013); Bourdieu (1993). A similar model concerning the development of status was introduced by Gould (2002).

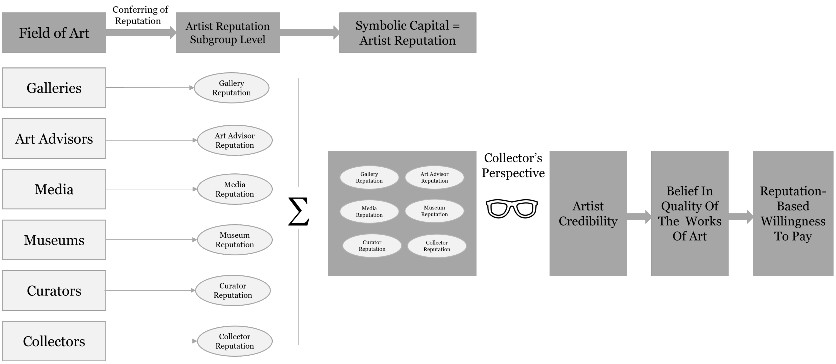

Figure 3.1: Model of the Credentialing Process

3.3 Methodology

3.3.1 Mapping the Global Art Community

Since its creation in 2010, Instagram has become the most popular social media platform in the art industry.[2] The platform enables users to share photos, videos, and other types of content through personal accounts that they curate themselves. Both individuals and institutions can create accounts on the platform. Users on Instagram can follow the accounts of other users, although the other users are not required to follow back. It is common for popular accounts with many followers to follow only a small number of other accounts. This study used data from 10,000 art-related Instagram accounts as a starting point for digitally mapping

[2]Hiscox online art trade report 2019: https://www.hiscox.co.uk/sites/uk/files/documents/2019-04/hiscox-online-art-trade-report-04-2019.pdf

the global art community network.[3] The Instagram accounts followed by the initial group of 10,000 art-related accounts may also be considered important within the art industry. Therefore, all the accounts followed by the initial group of 10,000 art-related accounts have also been included. Using this method, a total of 2,857,254 unique accounts were identified. The identified connections between accounts and their followers are used to create a network, with each identified Instagram account represented as a node in the network. In order to approximate the importance of the nodes (Instagram accounts) in the network, the indegree and eigenvector centrality measures from graph theory are used.[4] The top 4,095 nodes in the network, ranked according to indegree centrality, were manually classified into the following art-market subgroups: art galleries, art advisors, media, museums, curators, and others.[5] Although a node's relative centrality may change over time, it was assumed to be constant for the purposes of this analysis, because of the relatively short time period of observation. The top 100 Instagram accounts in terms of their indegree centrality statistics are listed in Appendix 3.A.

3.3.2 Measuring Status-Weighted Artist Attention

This paper uses Instagram as a source of data. On Instagram, users can upload pictures to their personal profiles and add a caption as a

[3]The initial account names were provided by NAMAC GmbH., a consulting firm for art market trends.

[4]For computational purposes, the network was pruned by keeping only the nodes that had indegree centrality above the 90th percentile. This resulted in a final network of 221,357 nodes.

[5]Of the top 4,095 representatives in the art market, as ranked by their indegree centrality in the network, 65% are also among the top 4,095 when the eigenvector centrality is used to rank the network.

description. A common approach to write the caption of a post is to use the hashtag symbol #. For example, someone might visit an art gallery to view a painting by Rashid Johnson, take a photo of it, and upload the photo to their Instagram account with the hashtag #rashidjohnson. This allows others who are interested in Rashid Johnson's work to discover the post and engage with it. Searching for #rashidjohnson will return all posts that have used the hashtag #rashidjohnson.

In order to measure the level of status-weighted attention an artist receives, the Instagram posts that include a hashtag created by combining the artist's first and last name are used. By requiring that the hashtag include both the artist's first and last name, the number of results is reduced, but it also helps to avoid false results if the artist has a last name that is very common. The value of a post on Instagram is equal to the node weight, or centrality, of the user who uploaded it. If the user who uploaded the post is not included in the art network community, the value of the post is considered to be zero.

Equation (3.1) describes the monthly attention value of artist i in month t by art-market subgroup k for $k \in$ {Galleries, Art Advisors, Media, Museums, Curators, Others}. $N_{i,t,k}$ represents the set of individuals or organizations within subgroup k who paid attention to artist i during month t. The status-weighted attention that artist i receives from subgroup k during month t is calculated as follows:

$$A_{i,t,k} = \sum_{j \in N_{i,t,k}} \alpha_j, \tag{3.1}$$

where α_j is the node weight, or centrality, of institution (person) j within subgroup k. The attention time series per artist were corrected for changing levels in overall Instagram activity over the years as in Serwart (2023).

3.3.3 Models

$A_{i,k,T}$ represents the average level of attention that artist i receives from art market subgroup k during year T.

$$A_{i,k,T} = \sum_{t=1}^{12} \frac{1}{12} A_{i,t,k}. \tag{3.2}$$

Ordinary least squares (OLS) regression is used to analyze the relative influence of an artist's reputation within different subgroups of the art market on their realized returns. The reputation of an artist is quantified using the status-weighted attention scores calculated using equation (3.2).

The dependent variable $r_{i,T}$ represents the return of artist i in year T. The return per artist $r_{i,T}$ reflects the change in price per square centimeter of the paintings produced by that artist. In order to calculate the return $r_{i,T}$ for an artist i in a particular year T, the ratio of the median price per square centimeter of the artist's paintings in year T

and year $T-1$ is taken. The return of artist i in year T can be stated as:

$$r_{i,T} = \frac{p^m_{i,T} - p^m_{i,T-1}}{p^m_{i,T-1}}, \tag{3.3}$$

where $p^m_{i,T}$ denotes the median price of artist i in year T. This return calculation methodology is only sensible if the artist has sold several paintings per year. Otherwise, the returns are estimated with high variance. In this paper, I require that an artist have sold at least 3 paintings in each subsequent year in order for the returns $r_{i,T}$ to be calculated. In the appendix, the results are regenerated for the restriction that at least 4 paintings have been sold in each year as a robustness test. This method of calculating returns using the median price per square centimeter was chosen because paintings do not often resell.[6] The analysis was limited to paintings to compare like with like as far as possible.[7] The model is:

$$r^l_{i,T} = \alpha + \sum_{k=1}^{K} \beta_{1,k}\, A_{i,k,T-1} + \sum_{k=1}^{K} \beta_{2,k}\, (A_{i,k,T-1} - A_{i,k,T-2}) + \sum_{i=0}^{1} \beta_i \text{ SP500 Return}_{T-i} + \epsilon_{i,T}, \tag{3.4}$$

[6]Another methodology to calculate artist returns would be to run hedonic regressions of the prices on artwork characteristics and year fixed effects per artist. The year fixed effects can be used to calculate annual returns as described in Renneboog and Spaenjers (2013). However, the small number of observations per artist and the comparatively large number of covariates that must be included in such a regression cause several problems. First, the model may not be estimable at all if there are more covariates than observations since the matrix cannot be inverted. Second, even if the model can be estimated, the standard errors of the year fixed effects are huge when the number of observations is approximately equal to the number of covariates. For these reasons, the methodology described in equation 3.3 was chosen as a second-best approach.

[7]Shortcomings of this methodology include the fact that the price per square centimeter of a painting is affected by numerous attributes such as its size, its age an its selling location, among others, as described in Serwart (2023). Therefore, the return calculations are not estimated net of these effects and are consequently likely to be subject to significant noise. For these reasons, a relative interpretation of the return figures across artists is more justified than an absolute one.

where $r^l_{i,T}$ represents logarithmic returns calculated as $log(r_{i,T}+1)$. The model regresses the returns on the lagged attention levels received from different subgroups within the art market. The model also regresses the returns on the difference in attention levels received from different subgroups within the art market between years $T-2$ and $T-1$. The time structure of this model helps to reduce the potential for reverse causality issues. The model also controls for the returns of the S&P 500 index, as increased liquidity in the market can lead to higher art prices when people's wealth increases. The attention values $A_{i,T,k}$ are calculated using either eigenvector or indegree centrality measures to determine the node weights of individuals or organizations within the art field.

The model suffers from reverse causality issues, which is why the coefficients cannot be interpreted causally. However, this is not a problem since this paper tries to develop an artist success prediction mechanism. For this purpose, causal interpretation of the coefficients is not required. The problem is illustrated in Figure 3.2. Specifically, the dependent return variable $r_{i,T}$ is calculated using the median price realizations. The median price realizations can take place at any point in time during any given year. However, the yearly averaged attention variables are calculated using attention scores of all months. Hence, for example the attention variable A_{T-1} includes the attention area $A1$, which is potentially affected by the return $r_{i,T}$, which is why the model suffers from reverse causality and the coefficients cannot be interpreted causally.

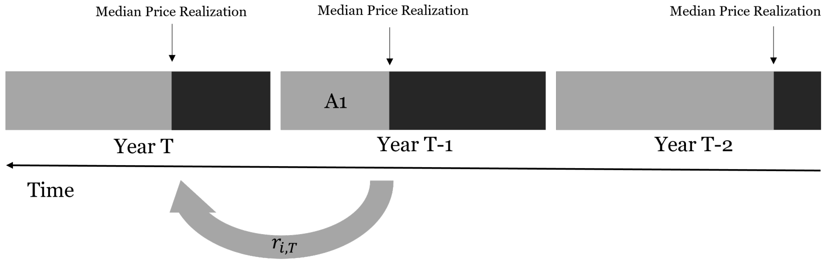

Figure 3.2: Prediction Explanation

The artist selection in the application exercise presented in Section 3.5 happens at the end of year $T-1$ and thus incorporates the attention information realized in area $A1$. This is still a valid prediction application, although it is simpler than a case where the prediction always takes place at the point in time of the median price realization.

3.4 Data

As in the study by Serwart (2023), this paper focuses on the return developments of emerging artists. There is no strict definition of the term *emerging artist.* In order to define this group, all artists who had sold at least one artwork in the *New Now* auction series at Phillips auction house were selected. The *New Now* auctions are held semi-annually in New York City and London and are dedicated to the sale of artwork by emerging contemporary artists. By using this selection criterion, it can be inferred that all of the selected artists have successfully sold at least one artwork at a well-known international auction house. Using this definition, the term *emerging artist* in this paper includes artists who have achieved some level of success and can be considered mature emerging artists. There were a total of 21 *New Now* auctions held between 2015 and 2021. In order to make the analysis more comparable, only artists who work with paintings were included in the study. Artists who work exclusively in other mediums such as sculpture or photography were excluded. This selection process resulted in a sample of 179 artists. For these artists, the study analyzed all artworks sold at all major auction houses, not just Phillips. The data used in the study were obtained from LiveArt. The final dataset includes 475 return observations for the 179 selected artists between 2015 and 2021. Table 3.1 provides descriptive statistics for the data. In Appendix 3.B descriptive statistics of the returns are shown when no restriction is imposed on the number of paintings sold in each consecutive year. The range of returns becomes wider with unreasonably high maximum returns. Furthermore, in Appendix 3.C descriptive statistics are given for the returns under the requirement of at least 4 sales during each consecutive year. The return numbers stabilize when more paintings are required to be sold in each consecutive year. However, this procedure also reduces the sample size. Therefore, the methodology here requires at least three paintings per year, while all results for the case of requiring 4 sales per year are given in the respective appendices.

	Total
Number of return observations	475
Number of distinct artists	179
Number of painting observations	5,293
Median hammer price ($)	37,647
Mean hammer price ($)	150,463
Maximum hammer price ($)	16,767,500
Minimum hammer price ($)	133
Observations in 2016	77
Observations in 2017	81
Observations in 2018	92
Observations in 2019	93
Observations in 2020	95
Observations in 2021	37
Minimum Return	−96.83%
10 Percentile Return	−58.02%
25 Percentile Return	−37.13%
50 Percentile Return	−1.99%
75 Percentile Return	45.89%
90 Percentile Return	118.51%
Maximum Return	1319.19%

Table 3.1: Descriptive Statistics of Return Data

This Table shows descriptive statistics of the return data. All prices were adjusted to real (2017) US Dollars. The returns were calculated under the restriction that at least 3 paintings were sold in each consecutive year.

The top 4,095 Instagram accounts in the various subgroups of the art market are summarized with descriptive statistics in Table 3.2. These subgroups encompass both individual and institutional players in the art world. In order to categorize individual Instagram accounts into the various subgroups of the art market, publicly available information was utilized. It is worth noting that a significant number of art collectors do not publicly disclose their collecting activity. In order to include these collectors in the analysis, a registry of contemporary art collectors from Larry's List was also used to identify their corresponding Instagram accounts.

Table 3.2 includes column 4, which displays the number of Instagram posts made by each subgroup in the sample for the 179 artists included in the data.

Art-Market Subgroup	***Number of Observations***	***Type of Accounts***	***Number of Artist Posts***
Galleries	1,020	Official Gallery Accounts, Gallery Owners	47,205
Curators	424	Known Professional Curators	9,525
Museums	390	Official Museum Accounts (Private and Public), Museum Directors	5,552
Collectors	337	Individual Art Collectors Official Institutional Collection Accounts	17,204
Media	262	Accounts of Art Magazines (Online and Offline) Newspapers, Journalists, TV Channels	6,902
Art Advisors	212	Art Advisors Art Advisories (Institutional Accounts)	13,904
Others	1,450	Offspaces, Art Foundations, Non-Profit Organisations, Artists Art Professors, Art Historians, Art Fairs	55,228
Total	4,095		136,808

Table 3.2: Art-Market Subgroup Descriptions
This table shows the different subgroups of art-market players and selected descriptive statistics of their Instagram profiles.

3.4.1 In-Sample and Out-of-Sample Data

In order to produce training and testing data sets, the data are split at the end of year 2018. The observations before and inclusive of 2018 are used to estimate model 3.4. The trading application, which is based on the results of model 3.4, is executed with data from years 2019-2021. This means 250 observations are used for training and the remaining 225 are used for testing in the trading application.

3.5 Results and Application

Table 3.3 shows that the difference in collector attention between $T-2$ and $T-1$ is the sole likely consistent predictor for artist returns between $T-1$ and T. In other words, if an artist experiences an increase in status-weighted collector attention between $T-2$ and $T-1$, ceteris paribus, one can expect his or her median painting price per square centimeter to increase between $T-1$ and T. As a robustness test, in Appendix 3.D in Table 3.9, model 3.4 is re-estimated requiring that at least 4 paintings be sold in each consecutive year. The results corroborate the hypothesis that the differential in collector attention is a useful return predictor. This is because the respective coefficients are the only ones to be statistically significant when the full model is estimated.

	Return per SQCM			
	Eigenvector	Indegree	Eigenvector	Indegree
Gallery -1	0.016	0.028		
	(0.084)	(0.077)		
Art Advisory -1	0.066	0.066		
	(0.100)	(0.097)		
Media -1	0.007	0.008		
	(0.116)	(0.114)		
Museum -1	−0.060	−0.061		
	(0.084)	(0.076)		
Curator -1	0.054	0.052		
	(0.119)	(0.112)		
Collector -1	0.018	0.040		
	(0.110)	(0.105)		
Other -1	−0.022	−0.047		
	(0.104)	(0.095)		
ΔGallery	−0.037	−0.046	−0.038	−0.045
	(0.079)	(0.072)	(0.069)	(0.061)
ΔArt Advisory	−0.130*	−0.118*	−0.090	−0.076
	(0.073)	(0.068)	(0.066)	(0.059)
ΔMedia	0.057	0.054	0.047	0.039
	(0.077)	(0.076)	(0.071)	(0.070)
ΔMuseum	−0.064	−0.042	−0.088**	−0.067*
	(0.060)	(0.053)	(0.043)	(0.039)
ΔCurator	−0.018	−0.049	0.005	−0.025
	(0.117)	(0.114)	(0.099)	(0.095)
ΔCollector	0.201**	0.190**	0.189**	0.172*
	(0.095)	(0.092)	(0.096)	(0.089)
ΔOther	0.004	0.025	−0.022	0.001
	(0.078)	(0.083)	(0.072)	(0.077)
SP500	−0.932	−0.936	−0.964	−0.980*
	(0.600)	(0.600)	(0.596)	(0.594)
SP500-1	−0.420	−0.451	−0.399	−0.431
	(0.743)	(0.749)	(0.752)	(0.759)
Constant	0.045	0.055	0.035	0.042
	(0.114)	(0.114)	(0.116)	(0.117)
Observations	250	250	250	250
Adjusted R^2	−0.002	−0.006	0.015	0.010

Table 3.3: Return Determinants Analysis

The table reports the OLS regression results of the annual returns per square centimeter per artist on different subgroup attention measures and controls. Arellano cluster-robust standard errors are shown in parentheses. The standard errors are clustered at the artist level. ***, ** and * denote significance at the 0.1%, 1% and 5% level, respectively. The variables names ending in a −1 reflect the attention measures for the specific subgroup in year $T-1$, while the variable names starting with Δ indicate the differential in the status-weighted attention measure of the specific subgroup between $T-2$ and $T-1$. The restriction that an artist needs to sell at least 3 paintings in each consecutive year is imposed.

This finding is exploited by developing the following application. First, in every year, a split variable is built, which describes the difference in collector attention between the two previous years $T-2$ and $T-1$. Those artists who experienced a large positive increase in weighted collector attention are more likely to have high positive returns between T and $T+1$. The realized returns of the top 10% (20%, 30%) of artists in terms of the split variable are compared to the realized returns of the bottom 10% (20, 30%) with respect to the split variable. The top and bottom group are always of equal size in terms of number of artists.

Table 3.4 depicts the results. Overall, the results do not depend on which network centrality measure was used to determine the node weights of the art world representatives. The results in Panel A and Panel B are very similar. It is apparent that the mean realized return of the artists in the top group over the years 2019-2021 is significantly higher compared to the realized return in the bottom group. The mean returns in the top group are between 60% to nearly 80% higher compared to the bottom group. The mean returns of the bottom and the top group are also statistically significantly different at least at the 5% level. Furthermore, the Sharpe ratios of the top group are economically higher and statistically significantly different from the Sharpe ratios of the bottom group. The minimum returns of the top and bottom group cannot meaningfully be distinguished. This also implies that there is not guarantee that the applied selection rule will mitigate large negative returns. However, the maximum returns are consistently higher in the top group. In other words, following the proposed selection rule helps to identify well-performing artists. Qualitatively, the same results obtain when requiring at least 4 sales in each consecutive year as depicted in Appendix 3.E in Table 3.10.

		Panel A: Eigenvector		
		10% / 90%	20% / 80%	30% / 70%
Top	Mean Return	65.1	71.9	60.2
	Sharpe Ratio	0.48	0.43	0.41
	Minimum Return (in %)	-67.1	-95.2	-95.2
	Maximum Return (in %)	439.46	780.26	780.26
Bottom	Mean Return	-9.3	-8	2.8
	Sharpe Ratio	-0.27	-0.23	0.01
	Minimum Return (in %)	-73	-94.9	-94.9
	Maximum Return (in %)	91.5	91.5	240.78
Comparison Statistics	T-Test p-value	0.013	0.002	0.003
	Sharpe ratio p-value	0.02	0	0.01
		Panel B: Indegree		
		10% / 90%	20% / 80%	30% / 70%
Top	Mean Return	55.4	67.2	61.8
	Sharpe Ratio	0.41	0.4	0.42
	Minimum Return (in %)	-67.1	-95.2	-95.2
	Maximum Return (in %)	439.46	780.26	780.26
Bottom	Mean Return	-13.5	-1.7	4.5
	Sharpe Ratio	-0.38	-0.07	0.04
	Minimum Return (in %)	-73	-75.8	-94.9
	Maximum Return (in %)	91.5	206.96	206.96
Comparison Statistics	T-Test p-value	0.021	0.008	0.003
	Sharpe ratio p-value	0.01	0.02	0.02

Table 3.4: Trading Application Top versus Bottom

This table shows the results of the trading application. The top (bottom) group depicts the results of the returns of the top (bottom) 10% (20%, 30%) artists with respect to the splitting variable describing the differential in weighted collector attention between $T-2$ and $T-1$. The comparison statistics test differences in the mean and the Sharpe ratio of the top versus the bottom group. Panel A (B) depicts the set of results for the case when *eigenvector* (*indegree*) centrality was used to determine the node weights (status) of the art representatives. The sample covers the years 2019 - 2021.

Moreover, a comparison between the top group and a random group is of great interest. This allows a judgment whether selecting artists based on the differential in past status-weighted collector attention generally yields superior returns compared to a random selection of artists. In every year, the top 10% (20%, 30%) of artists with respect to the split variable describing the differential in weighted collector attention between $T-2$ and $T-1$ are chosen and compared to a random subset of artists of equal size. The results of this return analysis are depicted in Table 3.5. Both the mean realized returns and the Sharpe ratios are significantly higher in the top group compared to the random group. The mean return is statistically significantly different in both groups at least at the 10% level when comparing the top 20% and 30% of artists against a random group.[8] The Sharpe ratio level differences are not statistically significant at traditional significance levels. The minimum returns in the top and random group cannot be distinguished meaningfully. This is not surprising since the minimum returns could also not be distinguished from the top versus bottom group in Table 3.4. However, the maximum returns are consistently higher in the top versus the random group.

In summary, this means that the proposed artist selection rule enables picking artists that generate on average economically and statistically significantly higher returns compared to a random selection of artists. While the selection procedure cannot guarantee that none of the chosen artists perform badly from an investment perspective, the selection procedure is very good at the early identification of subsequently very well

[8] The small sample size accounts for the fact that the mean return is not statistically significantly different in the top versus random group, when only the top 10% of artists are chosen and compared with a random selection of artists. The same reasoning applies to the finding that the mean return difference between the top and random group is not statistically significant when requiring at least 4 painting sales in each consecutive year, which majorly reduces the sample size. These findings are presented in Appendix 3.D in Table 3.11.

performing artists.

		Panel A: Eigenvector		
		10% / Random	20% / Random	30% / Random
Top	Mean Return	65.1	71.9	60.2
	Sharpe Ratio	0.48	0.43	0.41
	Minimum Return (in %)	-67.1	-95.2	-95.2
	Maximum Return (in %)	439.46	780.26	780.26
Random	Mean Return	24.8	19.4	23.9
	Sharpe Ratio	0.3	0.27	0.2
	Minimum Return (in %)	-67.4	-73	-73
	Maximum Return (in %)	191.71	191.71	780.26
Comparison Statistics	T-Test p-value	0.2	0.047	0.097
	Sharpe ratio p-value	0.48	0.33	0.1
		Panel B: Indegree		
		10% / Random	20% / Random	30% / Random
Top	Mean Return	55.4	67.2	61.8
	Sharpe Ratio	0.41	0.4	0.42
	Minimum Return (in %)	-67.1	-95.2	-95.2
	Maximum Return (in %)	439.46	780.26	780.26
Random	Mean Return	24.8	19.4	23.9
	Sharpe Ratio	0.3	0.27	0.2
	Minimum Return (in %)	-67.4	-73	-73
	Maximum Return (in %)	191.71	191.71	780.26
Comparison Statistics	T-Test p-value	0.333	0.068	0.088
	Sharpe ratio p-value	0.7	0.45	0.11

Table 3.5: Out-of-sample Random

This table shows the results of the trading application. The top group depicts the results of the returns of the top 10% (20%, 30%) artists with respect to the splitting variable describing the differential in weighted collector attention between $T-2$ and $T-1$. The results of the random group are below the top group results. The comparison statistics test differences in the mean and the Sharpe ratio of the top versus the bottom group. Panel A (B) depicts the set of results for the case when *eigenvector* (*indegree*) centrality was used to determine the node weights (status) of the art representatives. The sample covers the years 2019 - 2021.

For those interested in contemporary art, it is of interest to know which artists would have been selection candidates based on the proposed selection rule. Table 3.6 shows for each year T the names of those four artists with the highest differential in collector attention between $T-2$ and $T-1$. It would have been proposed to invest in these artists from an investment perspective at the end of year $T-1$. Four out of 16 artists show negative returns. The other 12 artists experienced very significant price increases per square centimeter, corroborating the observed efficacy of the selection rule.

	End-of-Year 2018 Selection	
	Name	***Return 2018-2019***
Top 4 Advise to Buy	Mark Grotjahn	-3.8%
	Nicole Eisenman	46.2%
	Sam Gilliam	20.8%
	Eddie Martinez	34.1%
	End-of-Year 2019 Selection	
	Name	***Return 2019-2020***
Top 4 Advise to Buy	Henry Taylor	-15.3%
	Nina Chanel Abney	128.0%
	Katherine Bradford	-41.0%
	Derek Fordjour	135.6%
	End-of-Year 2020 Selection	
	Name	***Return 2020-2021***
Top 4 Advise to Buy	Rashid Johnson	191.7%
	Katherine Bernhardt	-5.6%
	Dana Schutz	218.1%
	Eddie Martinez	439.4%

Table 3.6: Advise Illustration

This table shows which artists had the highest collector attention differential between the years $T-2$ and $T-1$ and would thus have been the artists most advised to buy at the end of year $T-1$. The return figures describe the (de)-appreciation in median price per square centimeter of the respective artist per paintings, between $T-1$ and T.

The flip side of this analysis is that it can also advise against buying certain artists in a given year. The four least-favored artists would have generated average returns of -18.4%, -37.1% and +15.08% in 2021, 2020 and 2019, respectively.

In summary, the analysis comparing the returns of the top group versus bottom group of artists in terms of their past weighted collector attention differential showed the following: From an investment perspective, avoiding buying certain artists with a relatively low collector attention differential compared to the top group allows generation of both economically and statistically significantly higher average returns. Furthermore, artist selection based on differences in past values of status-weighted collector attention also yields on average superior economic returns when compared to a random selection of artists. The differences are significant in economic and statistical terms. Only in one case, was the statistical difference in mean returns not statistically significant, which is likely due to the relatively small sample size in that case.

3.6 Conclusion

This paper applies a methodology that quantifies an artist's reputation over time. The novel Instagram data set offers the possibility of dynamically quantifying the reputation of an artist by tracking the status-weighted attention he or she receives from representatives of the field of art. The paper showed that past collector attention drives artist returns. Ceteris paribus, the higher the increase in past status-weighted collector attention, the higher the expected future returns of the artworks of that artist. In a trading application, it was shown that selecting artists based on their past track record of status-weighted collector attention dynamics entails several significant financial benefits. First, by avoiding investment in artworks of artists with a relatively lower collector attention differential in the past with respect to a comparison group of artists, significantly higher returns can be achieved in economic and statistical terms. Second, artist selection based on differences in past values of status-weighted collector attention yields on average superior economic returns compared to an art portfolio consisting of randomly chosen artists. The returns of the art portfolio constructed using the collector attention differential factor outperformed the randomly constructed portfolio by as much as 30% to 50%. The analysis should be of great interest to art investors and collectors alike, as it provides them with an easy-to-calculate quantitative measure to assist them in making decisions about the selection of artists from an investment perspective. Future research opportunities include among others applying and testing the developed methodology for non-contemporary art as well as constructing and applying quantitative reputation measures for producers in other fields that lack objective quality standards.

References

Beckert, J. (2020). Markets from meaning: quality uncertainty and the intersubjective construction of value. *Cambridge Journal of Economics*, *44*(1), 285-301.

Beckert, J., & Rössel, J. (2013). The price of art: Uncertainty and reputation in the art field. *European Societies*, *15*(2), 178-195.

Beggs, A., & Kathryn, G. (2009). Anchoring effects: Evidence from art auctions. *The American Economic Review*, *99*, 1027-1039.

Bourdieu, P. (1993). *The field of cultural production*. New York: Columbia University Press.

Goetzmann, W. N., Renneboog, L., & Spaenjers, C. (2011). Art and money. *American Economic Review*, *101*(3), 222-226.

Gould, R. V. (2002). The origins of status hierarchies: A formal theory and empirical test. *American Journal of Sociology*, *107*(5), 1143–1178.

Korteweg, A., Kräussl, R., & Verwijmeren, P. (2016). Does it pay to invest in art? A selection-corrected returns perspective. *The Review of Financial Studies*, *29*(4), 1007-1038.

Lovo, S., & Spaenjers, C. (2018). A model of trading in the art market. *American Economic Review*, *108*(3), 744-774.

Mei, J., & Moses, M. (2002). Art as an investment and the under-performance of masterpieces. *American Economic Review*, *92*(5), 1656-1668.

Pénasse, J., Renneboog, L., & Spaenjers, C. (2014). Sentiment and art prices. *Economics Letters*, *122*(3), 432-434.

Renneboog, L., & Spaenjers, C. (2013). Buying beauty: On prices and returns in the art market. *Management Science*, *59*(1), 36-53.

Serwart, J. (2023). Quantifying artist reputation: Disentangling art-market subgroup credentialing effects. *Poetics*. (Revise & Resubmit)

Appendix

3.A Top 100 Gallery and Museum Instagram Usernames

- 1. artbasel
- 2. artforum
- 3. tate
- 4. themuseumofmodernart
- 5. frieze_magazine
- 6. friezeartfair
- 7. artnet
- 8. guggenheim
- 9. davidzwirner
- 10. artsy
- 11. hauserwirth
- 12. hansulrichobrist
- 13. whitneymuseum
- 14. gagosian
- 15. momaps1
- 16. newmuseum
- 17. artnews
- 18. centrepompidou
- 19. metmuseum
- 20. mariangoodmangallery
- 21. whitecube
- 22. pacegallery
- 23. jerrysaltz
- 24. artreview_magazine
- 25. palaisdetokyo
- 26. moussemagazine
- 27. theartnewspaper.official
- 28. collecteurs
- 29. diaartfoundation
- 30. moca
- 31. serpentineuk
- 32. lisson_gallery
- 33. flashartmagazine
- 34. fiacparis

- 35. brooklynmuseum
- 36. lacma
- 37. artobserved
- 38. blouin_artinfo
- 39. klausbiesenbach
- 40. fondationbeyeler
- 41. sothebys
- 42. labiennale
- 43. thearmoryshow
- 44. thaddaeusropac
- 45. fondazioneprada
- 46. artinamerica
- 47. koeniggalerie
- 48. artspace
- 49. simondepury
- 50. christiesinc
- 51. thecollectorslist
- 52. galerieperrotin
- 53. whitechapelgallery
- 54. hammer_museum
- 55. victoriamirogallery
- 56. sfmoma
- 57. paulacoopergallery
- 58. sadiecoleshq
- 59. hyperallergic
- 60. alminerech
- 61. c.a.daily
- 62. guggenheim_venice
- 63. stuartcomer
- 64. levygorvy
- 65. saatchi_gallery
- 66. simonleegallery
- 67. spruethmagers
- 68. independent_hq
- 69. galleriesnow
- 70. xavierhufkens
- 71. petzelgallery
- 72. stedelijkmuseum
- 73. museelouvre
- 74. drawingcenter
- 75. anitazart
- 76. artbrussels
- 77. e_flux
- 78. britishmuseum
- 79. nespector
- 80. art.viewer
- 81. davidkordanskygallery

- 82. flagartfoundation
- 83. massimodecarlogallery
- 84. fondationlv
- 85. 303gallery
- 86. ceciliaalemani
- 87. museoguggenheim
- 88. vamuseum
- 89. icalondon
- 90. phillipsauction
- 91. walkerartcenter
- 92. the_adaa
- 93. gavinbrownsenterprise
- 94. rubellmuseum
- 95. bourgeois0490
- 96. newartdealers
- 97. kunstsammler
- 98. artissimafair
- 99. galeriemaxhetzler
- 100. larrys_list

3.B Return Descriptive Statistics - At Least 1 Painting

	Total
Number of return observations	1,221
Number of distinct artists	347
Observations in 2016	194
Observations in 2017	200
Observations in 2018	219
Observations in 2019	245
Observations in 2020	231
Observations in 2021	132
Minimum Return	-99.05%
10 Percentile Return	-66.08%
25 Percentile Return	-42.34%
50 Percentile Return	-5.52%
75 Percentile Return	54.82%
90 Percentile Return	183.8%
Maximum Return	$12,492.81\%$

Table 3.7: Descriptive Statistics of Return Data

This table shows descriptive statistics of the return data. These prices were adjusted to real (2017) US Dollars. This returns were calculated imposing no restrictions with respect to the number of paintings sold in each consecutive year.

3.C Return Descriptive Statistics - At Least 4 Paintings

	Total
Number of return observations	347
Number of distinct artists	133
Number of painting observations	$5,193$
Median hammer price (\$)	$38,194$
Mean hammer price (\$)	$152,509$
Maximum hammer price (\$)	$16,767,500$
Minimum hammer price (\$)	133
Observations in 2016	51
Observations in 2017	58
Observations in 2018	67
Observations in 2019	74
Observations in 2020	68
Observations in 2021	29
Minimum Return	-96.83%
10 Percentile Return	-56.99%
25 Percentile Return	-36.74%
50 Percentile Return	-1.68%
75 Percentile Return	43.77%
90 Percentile Return	128.71%
Maximum Return	899.57%

Table 3.8: Descriptive Statistics of Return Data

This table shows descriptive statistics of the return data. All prices were adjusted to real (2017) US Dollars. The returns were calculated under the restriction that at least 4 paintings were sold in each consecutive year.

3.D Robustness Test 4 Paintings per Year: Regression

	Return per SQCM			
	Eigenvector	Indegree	Eigenvector	Indegree
Gallery -1	0.003	0.0004		
	(0.107)	(0.100)		
Art Advisory -1	0.089	0.092		
	(0.138)	(0.133)		
Media -1	−0.181	−0.156		
	(0.178)	(0.171)		
Museum -1	0.022	0.011		
	(0.100)	(0.093)		
Curator -1	0.052	0.036		
	(0.127)	(0.119)		
Collector -1	0.077	0.089		
	(0.141)	(0.136)		
Other -1	0.010	−0.006		
	(0.104)	(0.097)		
ΔGallery	−0.041	−0.048	−0.059	−0.064
	(0.083)	(0.073)	(0.075)	(0.064)
ΔArt Advisory	−0.048	−0.041	0.028	0.032
	(0.101)	(0.093)	(0.085)	(0.076)
ΔMedia	0.130	0.118	0.054	0.044
	(0.096)	(0.092)	(0.073)	(0.072)
ΔMuseum	−0.011	0.019	0.001	0.024
	(0.076)	(0.069)	(0.053)	(0.049)
ΔCurator	0.001	−0.002	0.032	0.018
	(0.128)	(0.122)	(0.113)	(0.109)
ΔCollector	0.179*	0.160*	0.143	0.121
	(0.096)	(0.095)	(0.093)	(0.087)
ΔOther	0.011	0.031	−0.017	0.015
	(0.117)	(0.119)	(0.103)	(0.104)
SP500	−0.431	−0.462	−0.434	−0.461
	(0.716)	(0.710)	(0.705)	(0.701)
SP500-1	0.552	0.514	0.579	0.576
	(0.828)	(0.824)	(0.817)	(0.823)
Constant	−0.061	−0.053	−0.071	−0.067
	(0.122)	(0.123)	(0.122)	(0.123)
Observations	176	176	176	176
Adjusted R^2	−0.031	−0.033	−0.004	−0.005

Table 3.9: Return Determinants Analysis

The table reports the OLS regression results of the annual returns per square centimeter per artist on different subgroup attention measures and controls. Arellano cluster-robust standard errors are shown in parentheses. The standard errors are clustered at the artist level. ***, ** and * denote significance at the 0.1%, 1% and 5% level, respectively. The variable names ending in a −1 reflect the attention measures for the specific subgroup in year $T-1$, while the variable names starting with Δ indicate the differential in the status-weighted attention measure of the specific subgroup between $T-2$ and $T-1$. The restriction that an artist needed to have sold at least 4 paintings in each consecutive year is imposed.

3.E Robustness Test 4 Paintings per Year: Trading Application

		Panel A: Eigenvector		
		10% / 90%	20% / 80%	30% / 70%
Top	Mean Return	59.9	43.7	39.6
	Sharpe Ratio	0.47	0.36	0.37
	Minimum Return (in %)	-47.4	-95.2	-95.2
	Maximum Return (in %)	440.46	440.46	440.46
Bottom	Mean Return	-2.7	3.2	2.9
	Sharpe Ratio	-0.08	0.02	0.02
	Minimum Return (in %)	-71.6	-71.6	-72.5
	Maximum Return (in %)	115.57	206.96	206.96
Comparison Statistics	T-Test p-value	0.064	0.072	0.028
	Sharpe ratio p-value	0.09	0.09	0.03
		Panel B: Indegree		
		10% / 90%	20% / 80%	30% / 70%
Top	Mean Return	57.5	45.5	47.2
	Sharpe Ratio	0.44	0.39	0.41
	Minimum Return (in %)	-47.4	-95.2	-95.2
	Maximum Return (in %)	440.46	440.46	440.46
Bottom	Mean Return	-6.4	2.4	5.5
	Sharpe Ratio	-0.15	0.01	0.06
	Minimum Return (in %)	-71.6	-71.6	-72.5
	Maximum Return (in %)	115.57	206.96	206.96
Comparison Statistics	T-Test p-value	0.06	0.05	0.021
	Sharpe ratio p-value	0.05	0.07	0.04

Table 3.10: Trading Application Top versus Bottom

This table shows the results of the trading application. The top (bottom) group depicts the results of the returns of the top (bottom) 10% (20%, 30%) artists with respect to the splitting variable describing the differential in weighted collector attention between $T-2$ and $T-1$. The comparison statistics test differences in the mean and the Sharpe ratio of the top versus the bottom group. Panel A (B) depicts the set of results for the case when *eigenvector* (*indegree*) centrality was used to determine the node weights (status) of the art representatives. The sample covers the years 2019 - 2021.

		Panel A: Eigenvector		
		10% / Random	20% / Random	30% / Random
Top	Mean Return	59.9	43.7	39.6
	Sharpe Ratio	0.47	0.36	0.37
	Minimum Return (in %)	-47.4	-95.2	-95.2
	Maximum Return (in %)	440.46	440.46	440.46
Random	Mean Return	20.4	10.2	17.7
	Sharpe Ratio	0.26	0.12	0.2
	Minimum Return (in %)	-95.2	-95.2	-95.2
	Maximum Return (in %)	165.82	210.73	340.84
Comparison Statistics	T-Test p-value	0.25	0.15	0.223
	Sharpe ratio p-value	0.55	0.23	0.27
		Panel B: Indegree		
		10% / Random	20% / Random	30% / Random
Top	Mean Return	57.5	45.5	47.2
	Sharpe Ratio	0.44	0.39	0.41
	Minimum Return (in %)	-47.4	-95.2	-95.2
	Maximum Return (in %)	440.46	440.46	440.46
Random	Mean Return	20.4	10.2	17.7
	Sharpe Ratio	0.26	0.12	0.2
	Minimum Return (in %)	-95.2	-95.2	-95.2
	Maximum Return (in %)	165.82	210.73	340.84
Comparison Statistics	T-Test p-value	0.283	0.119	0.121
	Sharpe ratio p-value	0.6	0.19	0.18

Table 3.11: Out-of-sample Random

This table shows the results of the trading application. The top group depicts the results of the returns of the top 10% (20%, 30%) artists with respect to the splitting variable describing the differential in weighted collector attention between $T-2$ and $T-1$. The results of the random group are below the top group results. The comparison statistics test differences in the mean and the Sharpe ratio of the top versus the bottom group. Panel A (B) depicts the set of results for the case when *eigenvector* (*indegree*) centrality was used to determine the node weights (status) of the art representatives. The sample covers the years 2019 - 2021.

Curriculum Vitae

Education

Bachelor of Arts in Economics, University of St. Gallen, 2017.

Master of Science in Economics, London School of Economics and Political Science, 2018.

Ph.D. in Economics and Finance, University of St. Gallen, 2023.

Work experience

Analyst, c-alm AG, 2018–2020.

Research assistant, Faculty of Mathematics and Statistics, University of St. Gallen, 2020–2023.

Teaching

Data Analytics I: Statistics – Exercises, Bachelor in Economics, University of St. Gallen, 2020–2023.

Microeconomics – Exercises, Bachelor in Economics, University of St. Gallen, 2020.

Publications

Yield Curve Trading Strategies Exploiting Sentiment Data, *Revise and Resubmit at North American Journal of Economics and Finance*, F. Audrino, J. Serwart.

Quantifying Artist Reputation: Disentangling Art-Market Subgroup Credentialing Effects, *Revise and Resubmit at Poetics*, J. Serwart.