

A Site-Level Market Model of the Antiquities Trade

Fiona Greenland*

James V. Marrone†

Oya Topçuoğlu††

Tasha Vorderstrasse**

Abstract: Archaeological looting correlates with a number of problems, including the destruction of stratigraphic data and the damage and loss of artifacts. Looting is also understood to generate revenue, but systematic analysis of this issue is challenged by its opacity: how can we study the economic effects of archaeological looting when the practice is rarely directly observable? To address this problem, we estimate the market value

*University of Virginia Department of Sociology, Charlottesville, VA, United States; Email: fargreenland@virginia.edu

†RAND Corporation, Santa Monica, CA, United States

††Northwestern University, Evanston, IL, United States

**Oriental Institute, University of Chicago, United States

ACKNOWLEDGMENTS: The authors thank L. Brody for arranging access to the Yale University Art Gallery data, and Z. Dreyfuss, M. Vincent, T. Watler, and T. Wolter for research assistance. For helpful questions and comments on the paper, we thank N. Brodie, P. Gerstenblith, M. Gibson, M. Kersel, S.D. Kominers, B. Magness-Gardiner, G. Marsh, S. Parcak, L. Rothfield, G. Stein, M. Stephens, and the journal's anonymous peer reviewers. We are especially grateful to N. Judd and I. Hock for suggesting productive new methodological directions, and to participants in the Money, Markets, and Governance workshop in the Department of Sociology at the University of Chicago for perceptive comments and questions. Earlier versions of the paper were presented to audiences at the Neubauer Collegium, Department of Art History, and Oriental Institute at the University of Chicago; the Association for Public Policy and Management (APPAM); and the Annual Meeting of the American Institute for Conservation (AIC). The project was supported by grants from the Antiquities Coalition, the Oriental Institute at the University of Chicago, and the University of Chicago Department of Art History.

Artifact data records, constructed for this project from excavation reports, are stored in the OCHRE database at the Oriental Institute of the University of Chicago (<https://oi.uchicago.edu/research/ochre-data-service>). Market data, the data dictionary, and supplementary files are available online at <https://doi.org/10.1017/S0940739119000018>.

of archaeological sites where artifacts have been previously excavated and documented, using a machine-learning approach. The first step uses 41,587 sales of objects from 33 firms to train an algorithm to predict the distribution channel, lot packaging, and estimated sale price of objects based on their observable characteristics. The second step uses the trained algorithm to estimate the value of sites in which a large number of artifacts have been legally excavated and documented. We make an out-of-sample prediction on two Syrian sites, Tell Bi'a and Dura Europos.

Keywords: Archaeological looting, Syria, market valuation, machine-learning methods, Dura Europos, Tell Bi'a

INTRODUCTION

The trade in looted artifacts is a phenomenon of global proportions, adversely impacting community-level cultural knowledge and the scientific record and correlating with other forms of criminal activity.¹ While the trade is not new, recent conflicts in the Middle East have drawn intensified attention to endemic looting at archaeological sites in Syria and Iraq.² Quantifying the trade in monetary units has become an urgent research need because accurate estimates of money flows and profits hold promise for better policy instruments to counter the trade at both the source and in the marketplace.³ The interest of journalists and policymakers has been how much revenue particular groups generate by looting and selling artifacts. Scholars, meanwhile, have focused on archaeological site damage wrought by military operations, territorial expansion, and performative destruction, in addition to looting. The high-level revenue question and the fine-grained site study leave us with a considerable research gap. One major challenge is empirical observation; as with many illegal markets, direct empirical evidence is elusive.⁴ Interdictions capture an unknown fraction of illicitly exported artifacts and so must be used in conjunction with other data or certain strong assumptions in order to provide estimates of market size.⁵ Site observations (“ground truthing”) can reveal information about illicit digging frequency and intensity but cannot accurately capture artifact loss through digging. These problems lead to two related questions. First, on what basis can the site-level study be scaled up to region- or trans-national-level discussions about looting and market activity? Second, how can data on one phenomenon (in our case, scientific excavation) be made commensurable with data on another (artifact sales)?

¹See, e.g., Brodie and Renfrew 2005; Kersel 2006; MacKenzie and Green 2009.

²See papers in Casana 2015b; Casana and Laugier 2017.

³Fanusie and Joffe 2015; Parcak et al. 2016.

⁴Campbell 2013.

⁵The issues with such “supply-side” estimations of market size have been considered in other illegal markets as well. See Massey and Singer 1995; United Nations Office on Drugs and Crime 2005.

We approach this research problem at the level of a site. Prior studies of the antiquities market have used auction lots or pooled objects as the units of analysis. For our purposes—namely, studying a messy universe of artifact circulation and monetization—site-level analysis has important advantages, chief among them the availability of detailed excavation records that document the discovery and removal of artifacts from specific locations. In operationalizing our study this way, we acknowledge that the concept of “site” is a modern construct that delimits past flows of people, objects, practices, and ideas. Participants in the trade in archaeological materials, meanwhile, do not necessarily structure their work around particular sites. While some dealers and buyers specialize along lines of historical period, region, and artifact category, there is mixed evidence for site being a prioritized category in this market. We acknowledge these issues and agree that “site” should not be reified as the best or only way to approach the study of ancient societies. At the same time, the “site” as referring to a particular settlement or feature does have meaning in the study of looting because it is a practical referent for how illicit diggers organize their work.

One challenge to a site-level analysis is the great degree of heterogeneity among individual objects at a given site and between different sites from different periods of time. To this end, we rely on foundational work on archaeology and looting. Studies of archaeological looting have focused on damage to communities’ cultural heritage and loss of scientific data.⁶ This literature situated close interpretive analysis of specific sites and objects in broader socio-economic patterns to explain who looted and why and with what consequences for cultural knowledge and archaeological research. Subsequent studies explored quantitative approaches to clarify the scope and value of the trade,⁷ while scholars in law and criminology assessed the efficacy of legal instruments and enforcement mechanisms designed to stop it.⁸ What links these developments is methodology and scale. For the most part, scholars of looting have used qualitative methods to construct small-scale studies. There is great value in this approach as it elucidates subtle variations in structures and practices that ultimately account for macro-level pattern differences. A deficit, however, is that small-scale studies do not yield data that can explain looting and profit patterns at the regional, national, or global levels.

Recent developments in remote-sensing technologies have ushered in a new age of looting studies. Images from satellites and drones have allowed scholars to measure patterns in the frequency and extent of looting, covering large swathes of topography and tracking the development of looters’ pits with sophisticated metrics.⁹ Satellite and drone images do not, however, reveal what kinds of artifacts or how many were removed in the course of looting.¹⁰ Similarly, “ground truthing,”

⁶Coggins 1972; Vitelli 1979.

⁷Elia 2001; Kersel 2005.

⁸MacKenzie 2009.

⁹Contreras and Brodie 2010.

¹⁰Elia 2001; Casana and Panahipor 2014; Casana 2015b.

or on-the-ground examination of looted sites, can identify patterns documenting when and where illicit digging took place.¹¹ This technique is of limited assistance in estimating the market value of a site—by which we mean the sale price of ancient artifacts from the site, if they were dug up and sold at an established Western dealer or auction house—because it tends to focus on site and structure damage to the exclusion of quantitative market analysis.

Our work breaks new ground by combining market data with excavation data. Our unique data set includes 41,587 artifacts sold at 33 firms, consisting of 40 unique sales locations in seven countries, and 14,160 artifacts from two different excavation sites. Firm types include auction houses, brick-and-mortar private dealers, and online outlets. Methodologically, we classify each artifact's observable characteristics into variables and then use machine-learning versions of hedonic pricing to predict the object's sale location, sale format, and pre-auction estimate or posted price. We selected two sites—Tell Bi'a, an Early/Middle Bronze Age site in northern Syria, and Dura Europos, a Greco-Roman site in southeastern Syria—to make an out-of-sample prediction. Both sites were subject to systematic looting during the Syrian conflict.¹² Our model offers an important step toward estimating the market value of an entire archaeological site, thereby filling a major gap in the study of the illicit trade in artifacts. In order to estimate market value, we predict prices using a machine-learning algorithm. Millions of randomized predictions drawn from computerized simulations generate our estimates.

The core insight of our work is that identifying, aggregating, and matching physical characteristics of objects makes it possible to predict what sorts of items are buried at a site and how much they would be worth on the open market. The sites we selected for this study are representative of both their time periods and types of occupation (an Early/Middle Bronze Age regional administrative center and a Roman garrison), a fact that allows us to base our predictions on the assumption that the artifacts that have so far been found in the excavated portions of the site are representative of what could be uncovered at the rest of the site. Dura Europos, which was occupied over a very limited period of time, is particularly sensitive to this, as will be noted below.

We tackle the problem of how to study a largely unobservable practice by utilizing observable data in a new way. Our dataset includes small finds, a category of commonplace objects that includes, in the case of Greco-Roman sites, tools, jewelry, military regalia, and household goods and, in the case of Bronze Age sites, figurines, clay tablets, cylinder seals, household items, tools and weapons, jewelry, and decorative items. A key advantage of our data set, in comparison with previous scholarly studies, is its heterogeneity. We show that by working with a wide range of firms and artifact types, we can construct a model that is representative of the

¹¹Gibson 2003.

¹²Casana 2015b; *Illegal Excavations Are Still Going on at Mari and Dura Europos Archaeological Sites*, Damascus Government Museum Report, 27 February 2015.

diversity of the global antiquities trade. In turn, we generate insights into the huge, hidden market in antiquities that operates in the gray space between legal and illegal parameters.

DATA SELECTION

Two types of data were used in this analysis: objects offered for sale on the antiquities market and objects scientifically excavated and published.

Market Data

The market data consist of publicly available records of 41,587 sales of objects at a variety of international auction houses and dealers listed as members of the International Association of Dealers of Ancient Art (33 total firms). The pre-auction estimated or posted prices total \$469,116,026. Data consist of all sales information available online or in online catalogues for sales since 2007. Each object was atomized into fields describing five dimensions: period/date, size, constituent materials, type, and geographic/cultural source; in addition, market data have prices, which were converted to 2016 US dollars. These are the six observable characteristics that we consistently analyzed, in addition to the firm at which each object sold and the lot size in which it sold.

Our dataset offers several advantages based on the variety of firms we cover and the object-level analysis we are able to perform. We do not rely solely on data from top-tier auction houses. Catalogues from the most prestigious firms are an accessible source of information about the antiquities trade and have been subject to extensive scholarly study about provenance, market composition, and market size.¹³ As a data source, however, they present a distorted picture. Antiquities are sold by a variety of vendors and in a variety of groups, or “lots,” which sell for one collective price. Table 1 and Figure 1 show that the various combinations of vendors and sales format reveal differences in prices and object type. The table and figure group top-tier auction houses in New York and London (“large auction houses”), smaller auction houses, and private dealers. As Table 1 shows, large auction houses sell the most items overall, and they have the highest prices for items that are sold alone in their own lot. But large auction houses also sell the majority of their items through multi-item lots, unlike private dealers and other auctioneers. Large houses have the largest multi-item lots, on average, and the median price per item is actually the smallest.

These patterns are evident in aggregate in Figure 1, which shows that large auction houses have both the least and most expensive items, as evidenced by the lower and upper tails of the distributions. Small auction houses, on the other hand, have the smallest-size lots and the lowest prices for items sold alone. Private dealers sell most items individually, with prices higher than small auctioneers and around

¹³Chippindale and Gill 2000; Nørskov 2002; Brodie 2012.

Table 1. Characteristics of objects sold on market, by venue and lot size

	Large NY/London Auction Houses		Other Auction Houses		Private Dealers	
	Single-Object Lots	Multi-Object Lots	Single-Object Lots	Multi-Object Lots	Single-Object Lots	Multi-Object Lots
# Lots	11,998	20,214	1,624	756	4,488	2,507
Average # objects per lot	1	16.1	1	7.8	1	12.0
Median price-per-object (\$)	5,663	155	1,878	198	2,750	158
Average price-per-object (\$)	26,256	797	6,775	618	14,717	927
Median object size (cm)	18.0	6.5	13.5	7.0	11.4	8.0
Average object size (cm)	24.9	9.5	16.9	8.6	16.8	9.8
Bronze Age objects (%)	20.3	8.5	5.6	10.2	13.7	12.0
Greek objects (%)	22.1	12.6	7.4	7.4	18.9	10.2
Roman objects (%)	30.8	24.0	17.4	21.6	19.8	33.3




FIGURE 1. Market price distributions by venue: cumulative density functions representing price per object from the market data ($N = 41,587$) for various types of vendors.

50 percent of large auctioneers. In addition, various types of objects sell at different rates across venues. Greek, Roman, and Bronze Age artifacts (the types most probably found at our two test sites) make up the majority of objects at large auction houses, but the minority of objects at small ones. Bronze Age objects are relatively more likely to be sold at a private dealer, compared to Roman and Greek. And across all vendors, Roman objects are relatively more likely to be sold in groups rather than by themselves.

The price differentials between single- and multi-item lots distort the perceived average price and quality of the items that pass through top-tier houses.¹⁴ Because our analysis focuses on individual objects, not lots, our disaggregated data records the characteristics of each individual item. This is unusual in the literature, and our prediction model accounts for the lot-based sales practice of certain firms, allowing us to model outcomes more precisely. In particular, our algorithm calculates probabilities over various venue/lot-size combinations, based on the observable characteristics of each item. We can then calculate a distribution of prices for each object based on its predicted probability of being sold in a given setting.

¹⁴Marrone 2018.

While the data captured a wide swath of the market, sample selection is an important concern. By definition, we could not observe private sales or highest bids on items that did not meet reserve prices. Partly for this reason, we used pre-auction estimates rather than hammer prices. Overall, we believe pre-sale estimates and asking prices provide the best proxy for what market participants determine to be the *ex-ante* “expected” price of an object. In addition, previous research for other types of art suggests that, under reasonable assumptions, pre-sale estimates are unbiased predictors of the random final hammer price.¹⁵ We also could not observe any prices from the lowest end of the market; many small dealers do not post prices, and we deliberately excluded eBay sales because of concerns about authenticity and misattribution as well as the difficulty in viewing data from past auctions on the site.¹⁶ Inconsistency has also been a problem for other scholars working on antiquities auctioned on eBay.¹⁷ More generally, a concern regarding the use of such a diverse set of sellers is that the pricing strategies adopted by auction houses may be distinct from those adopted by private dealers due to the ways in which objects are offered for sale. However, this concern is alleviated by our computational approach, which trains machine-learning algorithms separately for large auction houses, small auction houses, and private dealers.

Excavation Data

We use the database of market sales to estimate prices for artifacts excavated from two archaeological sites in Syria: Tell Bi’a and Dura Europos. The test sites present two different chapters from Syria’s ancient history and, thus, represent different artifact assemblages. Looting has been widespread throughout the country, and sites from all of the major periods have been targeted. Tell Bi’a (Early/Middle Bronze Age) and Dura Europos (Greco-Roman period) differ from each other in material profiles, yet they share a common history of wartime looting. In addition, they represent two periods that are especially valued by collectors. Figure 2 shows the total value of objects in the market database by historical time period. Bronze Age, Hellenistic, and Roman objects are especially large market segments at the firms studied here, while segments such as Chalcolithic and Islamic are small. The finding of our algorithm corroborates on-the-ground observations that pre-Bronze Age sites yield objects that are of less interest to collectors; hence, they have experienced relatively low levels of looting.¹⁸ Thus, from a collector’s point of view, Dura Europos and Tell Bi’a represent high-value eras of Syria’s past.

In using these datasets for our model, we acknowledge that published excavation records have their limitations. As we highlight below, Dura Europos was excavated

¹⁵McAndrew, Smith, and Thompson 2012.

¹⁶For a discussion of authenticity and misattribution concerns, see Fay 2011.

¹⁷As pointed out over a decade ago by Kraft 2007.

¹⁸For key insights into this phenomenon in Iraq, see Stone 2008.




FIGURE 2. Market data by time period: aggregate market value of objects in the market database ($N = 41,587$) by century of origin. Dates correspond to the midpoint of an object's attributed range (for example, objects dated 400–200 BC are plotted at 300 BC or -3 on the graph). Periodization of the objects is based on sellers' descriptions/labels.

at a time when modern scientific excavation techniques had not fully developed, and not all aspects of the excavated structures and objects were accurately recorded. We also rely exclusively on the publications of the Yale University team, which only include the objects in their share, leaving out the material in the National Museum of Damascus. However, the Yale excavations covered approximately 40 percent of the entire site, so we are confident that our data are representative of the rest of the site. In the case of Tell Bi'a, careful and systematic excavation and recording methods allowed only 10 percent of the site to be excavated in 15 years. However, the site fits the general profile of Middle Bronze Age administrative centers in the region both in terms of its settlement pattern and artifact assemblage. As a result, the published data allow us to extend our predictions to the rest of the site.

In the case of Dura Europos, excavation data were drawn from the Yale University–French Academy team that worked at the site between 1928 and 1937, with online records totaling 12,398 objects; 11,566 of these are non-numismatic and were used for this study. After James Henry Breasted of the University of Chicago documented the wall paintings there, Franz Cumont led an excavation team in 1922–23. Cumont's work was followed by the joint Yale University–French Academy excavations, which lasted ten seasons. Following a long hiatus, the site

was excavated again from 1986 to 2011 by a Franco-Syrian team.¹⁹ The material from the Yale University–French Academy excavations was selected for detailed study here because it is the only one of the excavations to have produced a large number of detailed publications of the excavation, and the Yale University share of the excavation is available online, meaning that one has access to a large number of objects.

There are certain drawbacks to the corpus, however. Unlike scientific excavations today, the Dura Europos excavations did not carefully record stratigraphic levels or differentiated soil layers and the artifacts contained therein. Rather, they made general notes on the find-spots for the objects, with recording accuracy improving particularly from the fifth season (1931–32) onwards.²⁰ Nevertheless, patterns in the distribution of artifacts in specific parts of the site or across the entire site can be established because the site was built on a grid plan and is only occupied in the Greco-Roman period. The Yale records list each object with a size, cultural attribution, chronological date, physical description, and usually a photograph.

At Tell Bi'a, excavation data were drawn from reports published by the German team that worked at the site from 1980 through 1995. These sources include annual excavation reports that appeared in the *Mitteilungen der Deutschen Orient-Gesellschaft* periodicals between 1982 and 1995, six special volumes in the *Wissenschaftliche Veröffentlichung der Deutschen Orient-Gesellschaft* series dedicated to graves; finds from the Early Bronze Age levels; written documents; city walls, houses, and temples; and seals and seal impressions. The final book in this series, which covers the architecture and small finds excavated in Palace A, was published in 2007. As a recently conducted scientific excavation, Tell Bi'a offers the advantage of a complete dataset, in which all excavated objects were meticulously recorded, photographed, and published. These data constitute 2,594 total Bronze Age objects that were collected and entered into our project database by hand.

ARCHAEOLOGICAL BACKGROUND

Dura Europos

Dura Europos, a 72-hectare site in southeastern Syria, was founded at the end of the fourth century BC by the Greek Seleucids and was taken over by the Parthians in the late second century BC. After the middle of the second century AD, the Romans controlled the site until the Sasanian Persians conquered it in AD 256. There is limited evidence for occupation before and after this general timeline.²¹ As such, the site offers clear start and end dates of human inhabitation, allowing us to develop our predictive methodology on a specific subset of objects. Two more

¹⁹Comprehensive discussion of the data is offered by Baird 2012, 2014; Brody 2011.

²⁰Baird 2012.

²¹Baird 2012, 2014; James 2012.

factors make Dura Europos a good case study: it is associated with object types that have high collecting interest (as shown by Figure 2), and its geomorphology offers favorable conditions for preserving a wide range of artifacts including natural fibers, wood, and leather, thereby providing a diverse dataset for analysis. However, these factors also make Dura Europos a unique case among Greco-Roman sites in Syria in terms of its unusual artifact assemblage and objects that could be excavated and placed on the market.

Coins from Dura Europos have not been considered in this analysis for several reasons, despite the fact that they are primary targets of looting. While coins are, of course, a type of archaeological artifact, when it comes to the modern antiquities market, they are treated as a different type of good in terms of how they are marketed at auction and how their value is determined. Coins are usually sold at specialized auction houses and, therefore, largely do not overlap with the auction houses that were studied here. Indeed, coins constitute only 35 of the 41,587 items in our database, most having been made into jewelry in the modern period. From the point of view of our algorithm, the main difference is that coins have idiosyncratic characteristics that are important in determining price: mint, for example, is an attribute that only coins have. The value of coins is important to our understanding of the scale of looting in Syria and will be considered by Oya Topçuoğlu and Tasha Vorderstrasse in a separate, forthcoming publication on small finds.²²

Reports indicate that heavy looting had already occurred at Dura Europos by mid-June 2013, when the on-site museum was damaged and several large-scale items were taken from the site without official permission.²³ Dura Europos came under the control of the Islamic State in Iraq and Syria (ISIS) in September 2014, at which point it had sustained “large-scale, systematic looting” that had “destroy[ed] the archaeological record over an enormous area.”²⁴ After an apparent pause in looting activity from autumn 2014, Dura Europos was again subject to illegal excavations at least by April 2015.²⁵

Tell Bi'a

Tell Bi'a (ancient Tuttul), a Bronze Age mound, is a 52-hectare site located in the Ar-Raqqa governorate in northern Syria, 800 meters north of the city of Raqqa, once the headquarters of ISIS. The site was excavated between 1980 and 1995 by a German team led by Eva Strommenger. The settlement was first occupied in the third millennium BC. This period of occupation dating to the Early Bronze Age

²²Topçuoğlu and Vorderstrasse forthcoming.

²³See, e.g., papers in Casana 2015b.

²⁴The quotations are from Danti 2014. Samuel Hardy, “Was Most of the Looting at the Most Extensively Looted Site in Syria Conducted under the Islamic State or the Free Syrian Army?” *Conflict Antiquities*, blog, 12 March 2015, <https://conflictantiquities.wordpress.com/2015/03/12/syria-islamic-state-free-syrian-army-paramilitary-funding-dura-europos/>.

²⁵Casana and Laugier 2017.

yielded a large section of the city wall as well as the remains of a palace (Palace B), a temple (Temple C), and houses and graves, located in the southwestern quadrant of the settlement. Excavation reports also mention an as yet unexcavated palace (Palace C) and royal graves dating to the Early Bronze Age. The city reached its height during the second quarter of the second millennium BC (circa 1790–76 BC), when it functioned as one of the administrative centers of the Kingdom of Upper Mesopotamia ruled by king Samsi-Adad I. This period of occupation is best represented by a large palace (Palace A) at the center of the settlement, which was the focus of the German excavations. The building yielded large numbers of cuneiform tablets, seal impressions, and other small finds, such as ceramics, gaming boards, figurines, and clay models of houses and wagons. After its conquest and destruction in 1776 BC by King Zimri-Lim of Mari, Tell Bi'a appears to have lost its importance as a regional center. Late Bronze Age levels at the site have only yielded the remains of private houses. The site went through a revival during the Byzantine and early Islamic periods, which are represented by a sixth century AD Christian monastery and an Islamic coin hoard. However, this later inhabitation at the site is not considered in our analysis because it has not been fully published and other comparative Bronze Age sites do not all have similar reoccupation.

Satellite imagery of Tell Bi'a from as early as October 2012 shows dense looting in the form of hundreds of looters' pits, concentrated in the southwestern quarter of the site around several Early Bronze Age structures and in the area around the Early/Middle Bronze Age palace (Palace A), where a high concentration of small finds was excavated.²⁶ In publicly available satellite image of the site from February 2016, provided by Google Earth/DigitalGlobe, smaller concentrations of pits are also visible in other parts of the site. In terms of its archaeology and artifact assemblage, Tell Bi'a is representative of Bronze Age administrative centers across northern Syria. A number of Bronze Age sites comparable to Tell Bi'a in terms of size and relative importance (that is, major royal capitals and administrative centers) are found within the areas that experienced the heaviest wartime looting. The analysis of excavated objects from this site as well as satellite imagery showing evidence of large-scale industrial looting can inform future analysis of these other sites, which have either been already looted or are at risk of being looted.

STUDY DESIGN

We developed an original framework using standard machine-learning tools to identify site-level market values that can form the foundation of a more comprehensive analysis of region-wide market values. Our premise was that revenues from the trade in looted artifacts are best estimated not by scaling up from disparate anecdotal stories about a single object—a common practice in media stories and some scholarly publications—but, rather, via full-site data aggregation and hedonic

²⁶Casana and Panahipor 2014.

analysis of object characteristics. Based on our dataset, we would expect that site-level estimation would yield a more representative range of prices than that found in auction-only studies. If this is the case, then the overall expectation is that one-off high-value artifact sales are overrepresented in crude estimates. This does indeed appear to be true in our data; as discussed further below, objects valued over \$1 million account for only 0.05 percent of the objects but 16 percent of the market value. One item, the famed and record-breaking Guennol Lionness, accounts for 3.4 percent of the market on its own. In fact, the majority of objects recovered from sites comprise mundane objects with modest market values,²⁷ and our itemized dataset contains a greater proportion of such objects than does any previous sample.

To account for this observed distribution of objects and prices, any price-prediction algorithm needs to address outside influence of extremely valuable, but extremely rare, objects. Our analytical approach does so by using a large number of descriptive variables that optimize the fit for the vast majority of objects while sacrificing a perfect fit in the uppermost and lowermost tails. The conceptual approach differs from some previous studies of the illicit artifact trade; rather than treat a given artifact as a singular actor in the antiquities trade, we atomize the artifacts in our data set into five characteristics: time period, type, size, composite materials, and geographic origin. Breaking down artifacts based on observable characteristics is common practice in any art market, at a minimum as a way of appraising objects to calculate reserve and asking prices. It is also common practice in economic studies of the art market, which use hedonic regressions to analyze the importance of different characteristics in determining price.²⁸ Similarly, archaeologists classify artifacts according to period, object type, physical dimensions, composite materials, find spot, and other characteristics depending on the researcher's needs. But while specialists in the trade and in archaeology routinely sort by sub-characteristics, studies of the antiquities trade have not done this in any systematic statistical fashion. Our database integrates both archaeological data and market data, and, as a result, we are able to apply statistical methods to the antiquities trade in a way not previously possible. We have found that hedonic pricing techniques sustain more robust calculations of valuations both within object categories and across types.

Our prediction algorithms rely on tree-based machine-learning methods, which improve on standard hedonic regression techniques by finding the optimal set of interactions between our large numbers of descriptive variables. We use two different algorithms: random forests and Bayesian additive regression trees (BART).²⁹ Both algorithms are discussed in more detail in Appendix 1. As part of the price prediction algorithm, we introduce a novel innovation to the standard hedonic pricing approach by simulating a firm/lot size combination for each item, using

²⁷A point also made by Elkins 2012, 7.

²⁸Chanel, Gérard-Varet, and Ginsburgh 1996; Ashenfelter and Graddy 2006.

²⁹On random forests, see Breiman 2001; for Bayesian additive regression trees (BART), see Chipman, George, and McCulloch 2010.

predictions based on observable characteristics. This method accounts for important, but as-yet-overlooked, aspects of market structure and segmentation. As a result of this algorithmic approach, we can supplement our market valuation estimates with sub-analyses of specific object types, allowing us to address questions about time to market, market segmentation between firms, final sale context, and aggregate values of small finds.

Market Segmentation

Before calculating site-level valuations, our variables provided enough information to predict where and how an object will be sold, and we identified some surprising aspects of market composition. Table 1 and Figures 1 and 2 highlight certain salient aspects of market segmentation. These and other market characteristics highlighted here inform the way in which we performed the valuation calculations.

Object size is another characteristic by which our market data offer new insights. Using the 41,587 market data observations, Figure 3 shows the total market value by object size (measured as the object's largest dimension). Previous studies of the art market find that size is an important predictor for price: conditional on other characteristics, people tend to pay more for larger objects.³⁰ The mean object size in the raw data is 15 centimeters, with a median of 10.2 centimeters. But Figure 3 shows that objects in the range of 5–10 centimeters have the largest market share, with objects under 30 centimeters (roughly one foot) accounting for 88 percent of the objects and 71 percent of the market value. This finding is consistent with previous studies that demonstrate that the most ubiquitous items are small enough to be carried in a looter's pocket; our data add to this picture by showing that the market value is due to sheer numbers, not to value. For example, objects under 30 centimeters are priced at about 20 percent less than the average for the whole sample and 82 percent less than objects over 30 centimeters.

As for price, the data show that despite vast amounts of heterogeneity in objects' idiosyncratic characteristics, 95 percent of them have pre-sale estimates below \$32,000 (see Figure 1). In fact, only 0.05 percent of objects have estimates above \$1 million. This reveals a market that is more homogeneous than might be inferred from news reports, a reassuring fact given that our set of covariates simply cannot account for every possible characteristic of each object.

Overall, the market and archaeological data together demonstrate that the market is replete with small objects with no discernible aesthetic or historic assets for collectors—for instance, incomplete and corroded pieces of metalwork, including scale armor and other pieces of military equipment. Despite this fact, the artifacts are collected by individuals, and, therefore, it is critical to consider them. Small finds such as these also have clear advantages. They are easy to smuggle; what is more, our analysis predicts that those objects that are not worth much money will likely be aggregated into

³⁰As discussed by Beggs and Graddy 2009.




FIGURE 3. Market data by size. The histogram shows the aggregate market value of all items in the market database, grouped by largest dimension in five-centimeter increments.

large lots to achieve a high total price. These ubiquitous objects, however, have been largely overlooked by policymakers and the press, for whom unique, monumental, and very expensive artifacts are the chief concern. Our analysis reaffirms the observation made by David Ilan, Uzi Dahari, and Gideon Avni nearly 30 years ago: small finds are the engine drivers of the antiquities black market.³¹

Site-Level Estimates

Aggregating over the simulated sales outcomes based on a BART model calibrated with the market data, the mean value of the objects excavated at Dura Europos is \$18,088,000 with a 95 percent confidence interval of \$17,445,000 to \$18,765,000. At Tell Bi'a, the estimate is \$4,064,000 with a 95 percent confidence interval of \$3,883,000 to \$4,268,000. What accounts for the difference between sites? Dura Europos is larger (72 hectares versus 52 hectares at Tell Bi'a), and a higher percentage of the site was excavated (approximately 40 percent versus less than 10 percent at Tell Bi'a) and, as a result, yielded more artifacts, as reflected in our data set, which includes 11,566 artifacts from Dura Europos versus 2,594 from Tell Bi'a. In addition, Dura Europos was excavated in the 1920s and 1930s,

³¹Ilan, Dahari, and Avni 1989.

at a time when archaeology prioritized the recovery of artifacts over the recording of complete stratigraphic information, while Tell Bi'a was excavated some 60 years later, when excavation techniques had developed considerably. This has implications for how much of the site was excavated and the number of artifacts recovered. At Tell Bi'a, by contrast, the excavation work proceeded with slow, careful recording of stratigraphic layers. This systematic method also explains why we have more detailed documentation of the site. Finally, Dura Europos and Tell Bi'a were occupied 1,200 years apart and served different political and civic purposes; Dura Europos being a Roman garrison and Tell Bi'a being a regional administrative center. As such, they present different artifact profiles, which are of interest to different types of collectors.

How do we know whether these estimates are valid? Validating the model fit on out-of-sample market data provides support for the primary assumption underlying our methodology—namely, that the extrapolated prices for excavated objects are valid, insofar as the excavated objects share the same characteristics as items on the market, and that the unexcavated areas contain similar materials if extrapolating out to the whole site. It is possible that our excavated sample contains unique items that, were they to come to market, would be sold at high enough prices to dwarf our estimates. However, even if this were true for a handful of items, our estimates would still reflect the value of the vast majority of objects. Again, the market data provides a benchmark. Figure 3 shows that small objects comprise most of the market value; Figure 4 shows this is also predicted to be true for the excavated objects. At Dura Europos and Tell Bi'a, respectively, 50 percent of the predicted market value is attributed to objects under 13 centimeters and 7 centimeters, and upwards of 80 percent for objects under 30 centimeters.

In addition, the market segmentation discussed above allows us to verify that the algorithm makes realistic predictions regarding the “where and how” of hypothetical market sales for excavated items. If antiquities were randomly dispersed and sold in random combinations, then these variables would have no predictive




FIGURE 4: Predicted market value of excavation data, by size. The histograms show the predicted aggregate market value of all items from Dura Europos (left) and Tel Bi'a (right), grouped by largest dimension in five-centimeter increments.

power in determining price. Table 1 implies that the sales venue and lot sizes are not random. Our hedonic pricing model (Step 2 in Appendix 1) tests whether this is true by using the market dataset to calibrate a prediction algorithm. The algorithm then calculates the probability that a given excavated object sells at a given firm in a lot of a given size. Figure 5 shows the results of this prediction: each line corresponds to an object in the archaeological dataset, with color indicating the estimated average price of the object across all possible venue/lot size combinations. The height of the line varies to describe the probability that the object sells at a given firm or in a given lot size. The dark black lines show what the probabilities would be if an object were to sell randomly according to the empirical distribution of objects in the market dataset.

Figure 5 reveals two main points. First, Dura Europos contained the majority of higher-value items (that is, those valued in thousands of dollars rather than hundreds). Second, items are not distributed randomly according to the distribution of items in the market as a whole (if they were, all of the lines would overlap with the dark black lines). Instead, each item is more likely to sell in some contexts than in others. In particular, most items are more likely to sell at two particular top-tier auction houses or at a private dealer rather than at the other auction houses.

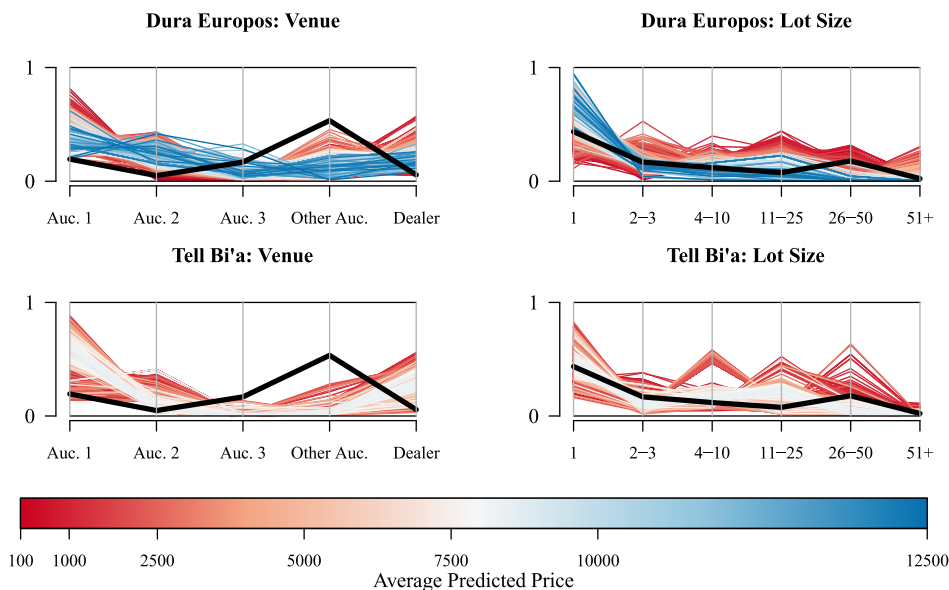


FIGURE 5. Predicting sales venue, lot size, and price. Each line represents an object in the market database. Line height corresponds to the probability the object is sold at a given venue or in a given lot size. Venues indicate the top three auction houses individually, other auction houses collectively, and private dealers. Line color indicates average expected price, accounting for the predicted distribution of all venue/lot size combinations. For clarity, scale is nonlinear and the most expensive 1 percent of items were dropped from the figure.

Thus, some sellers specialize in the types of Near Eastern objects found at Dura Europos and Tell Bi'a, relative to the rest of the market. In addition, items with a high estimated value are also more likely to be sold alone or in small lots, while those with low estimated value are more likely to sell in medium- to large-sized lots, relative to the market as a whole. This reflects the convention of grouping low-value objects into large lots and selling them for a single combined price.

The predictions plotted by Figure 5 are based only on the observed characteristics of each item—in other words, those characteristics that are known at the time of excavation. Yet the figure implies that firm and lot size not only segment the market based on these characteristics but also correlate with price. Overall, Figure 5 yields lessons for future research: different items have different predicted propensities to wind up in different sale contexts, and the determinants of the final point of sale must be considered as an important factor in reconstructing looting networks and supply chains.

In emphasizing lower-valued objects that make up the majority of market value, our approach is supported by a growing literature that emphasizes small finds and homogeneous objects as the focus of scholarly inquiry. To further validate our method, other machine-learning algorithms could be used, and alternative assumptions could be made. We found that the estimates were robust to alternative sets of descriptive variables used in the hedonic pricing algorithm—for instance, when using a coarser set of variables that ignored detailed descriptions of objects' shape, details, and surface designs.

DISCUSSION AND FUTURE DIRECTIONS

The above figures are our estimates of what the corpus of recorded, excavated artifacts from Dura Europos and Tell Bi'a would fetch on the market in 2016 dollars. To calculate the market value of looted objects from these sites, we are limited to parsimonious calculations. Nobody knows exactly what has been unearthed from looters' pits. At Dura Europos, for example, Yale University excavated about 40 percent of the city and took about half of its excavated artifacts, giving the remaining share to museums in Damascus and Paris. When the Yale team arrived in the 1930s, there were already reports of looting at the necropolis.³² Looting continued in the 1980s and 1990s; according to satellite data, the area within the city walls was hit hardest at this time. Clearly, any calculation of what may have been earned during the current conflict must account for this record of looting.

Our out-of-sample estimates are not, on their own, sufficient to make determinations about illicit monetary flows to market actors or non-state groups. They should be seen as an important first step. To undertake the next step, additional

³²Cumont 1925, 251; Toll 1946, 2–3.

information is necessary: first, which areas were looted in which time periods and, second, the payment structure between looters and secondary market makers.³³ To take the first point, Dura Europos had a long history of looting prior to the outbreak of the Syrian conflict in 2011.³⁴ Second, the final price does not equal what the looters or their controlling middlemen will have pocketed. Prior work on the illicit artifacts market suggests that looters and middlemen pocket only a fraction of the final price. In one study, Neil Brodie found that looters earned between 0.7 percent and 1.4 percent of the final price.³⁵ Building on Brodie's discussion, our data and analysis allow for further understanding of black market cash flows.

Other studies have established farm gate prices that can be compared with categories of objects in our database. In her study of southern Iraqi sites in 2007, Joanne Farchakh-Bajjalay found that "a cylinder seal or a cuneiform tablet earns \$50" for a looter.³⁶ Based on our data, this equals 2–7 percent of the average and median prices, respectively, of all cylinder seals. From interviews with looters in northern Jordan, Jerome Rose and Delores Burke found that they are paid \$15 per Roman or Byzantine bracelet, which we calculate as only 0.2–0.3 percent of the final market price.³⁷ (In the case of ISIS, it was widely reported that the regime extracts a cut of the projected market price on the order of 20–50 percent.³⁸) Additionally, there is currently insufficient evidence for differential cuts of profit among different participants in the trading network. Our analysis is a key step towards a more complete understanding of the cash flow structure.

The site-level market value estimate represents a significant advancement in the study of the circulation of archaeological artifacts in the international antiquities market. In one sense, our model is a mathematical relationship that can be used to predict prices. Policymakers and international non-governmental organizations have struggled to estimate market values, in part because they lack reliable data and robust models. Our empirical approach and model offer a solution to the problem of reliability and robustness. Moreover, the model itself has several possible functions. The first concerns the market for antiquities. Addressing the immediate problem of how to study the economic effects of archaeological looting when the details of the practice are not directly observable, we combined publicly available market and excavation data and applied tried-and-tested economic tools in a novel way. We demonstrated that it is possible to improve on previous research, which was generally object focused and had little to no explanatory power for site- or even region-wide looting observations.

³³Campbell 2013.

³⁴Casana 2015a, 147.

³⁵Brodie 1998.

³⁶Farchakh-Bajjalay 2008, 50.

³⁷Rose and Burke 2014.

³⁸Fanusie and Joffe 2015, 9–10.

To this end, our results reveal three important features of the trade in looted antiquities. First, we have added substantive findings to the emergent work on the trade in small finds. Small, commonplace, and mundane artifacts receive relatively little attention from scholars and policymakers concerned with archaeological looting, and, yet, they are massively present at auction, and their modest prices add up to a sizeable market.³⁹ Second, we have isolated the characteristics that make certain antiquities more desirable on the market. Objects from the Greek and Roman periods are in high demand, as are pieces made of gold, bronze, or glass and figurines. Prior to our analysis, experts suspected that these objects were favored on the market; we verify this suspicion statistically and provide clear categorical distinctions. Finally, systematic comparison between items on the market with the excavated artifacts from two sites has taught us that nearly every item, no matter how small, can find a buyer and that market participants may be dividing up significant numbers of marketable objects from just a single site. Continued refinement and application of our model will undoubtedly teach us much more about the function of archaeological sites in the global antiquities trade.

More broadly, our model offers a powerful way of thinking about excavation and market data. The site model can be used as a component of larger modeling systems for country-level or global-region political economy topics. We see our model as a first step toward a much larger research agenda, one extending into the economic development and resource conflict literatures.⁴⁰ Scholars have had little to say about artifacts as a contested, non-renewable resource, in part because artifacts' macroeconomic indicators are poorly understood. The site-level market model offers more analytical flexibility than gross domestic product or other nation-level indicators of resource revenue. It is too early to say whether artifact deposits are a "curse" or a "blessing" for a country or whether they are correlated with the duration of civil wars.⁴¹ What we can say is that archaeological sites present non-incident market values as well as opportunities for looters and smugglers to participate in heterogeneous streams of illicit goods. Moreover, the model can support the work of scholars and cultural resource managers concerned about protecting archaeological sites from exploitation and damage—whether from looting, agriculture, industrialization, or climate change. The ability to quantify cultural and scientific loss presents a building block toward clearer policy objectives. The outcome of such analysis does not have to be market value; our model allows for robust predictions of object quantities and types. Archaeological artifacts can no longer be thought of as ancillary goods in global development and conflict. Machine-learning methods using systematic data have the potential to open a new horizon of analysis of the social, scientific, and market costs of looting.

³⁹Topçuoğlu and Vorderstrasse forthcoming.

⁴⁰Ross 2006.

⁴¹Haber and Menaldo 2011; Fearon 2004.

APPENDIX 1

In this appendix, we explicate the steps for estimating the market value (in 2016 US dollars) of artifacts already excavated and recorded. The core of the method is a random-tree machine-learning version of the hedonic pricing techniques commonly used by economists to characterize auction prices and bidder behavior.

Step 1: Data Dictionary/Data Preparation

The nature of the antiquities trade presents two idiosyncrasies when trying to make different sources of data commensurable. The first idiosyncrasy is that multiple objects are often sold together in one “lot.” The second is that different sellers use different conventions to describe similar objects. To deal with the former problem, we used information on the numbers and details of objects sold together. In cases where multiple objects are sold together, we generated a separate observation for each object and divide the total price by the number of objects. We omitted cases where the number of objects in a lot could not be determined.

To handle the second problem, we generated a standardized corpus of descriptive terms and recoded objects’ descriptions based on this corpus. This is particularly important for price imputation; objects must be deemed “comparable” by some objective standards, while sellers maintain heterogeneous conventions for describing objects. We cannot be sure how any excavated object will eventually be described, if it is publicly sold at all; hence, we collapsed the hundreds of unique object descriptors into a set of 320 words indicating cultures, physical material, time period, and other descriptors. From this standardized corpus, each object could be described along the five main dimensions: size, time period, culture or geographic region, physical material, and type. The particular words in the corpus and the recoding strategy form part of our methodological contribution. The Data Dictionary in the supplementary materials file provides complete directions for constructing the standardized descriptive variables.

Step 2: Calibration

Next, we used a tree-based machine-learning approach to generate a predictive model of price. Tree-based classifiers are a natural choice because most of our variables are binary, which lend themselves nicely to categorization into sets of objects with the same combinations of characteristics and because we expect the optimal prediction model to have a hierarchical structure since certain characteristics will be relevant determinants of price only when other characteristics are present.⁴²

⁴²Hastie, Tibshirani, and Friedman 2009

To leverage these characteristics of the data, we used random forests and BART to develop prediction models without imposing assumptions regarding which variables are important. Of the full market dataset, 12,154 observations had missing information on size and/or date of origin. These missing values were imputed using random forests.⁴³ This and every subsequent implementation of random forests used 2,500 trees with 50 variables per node, although the results are not sensitive to parameter values above 1,000 trees and 30 variables per node.

Next, to aid in the prediction of where and how an excavated item will be sold, random forests calibrated a prediction of sales venue and lot size based on each object's observable characteristics. For this purpose, the 33 firms were split into five groups: a separate group for each of the top three auction houses, one group comprising all other auction houses, and a group comprising private dealers. Lot sizes were split into six groups: 1 item, 2–3 items, 4–10 items, 11–25 items, 26–50 items, and 51 or more items. These classifications were developed to maximize the accuracy of the prediction algorithm while maintaining a reasonable number of categorical groups. Random forests were then used to predict




FIGURE 6. BART convergence diagnostics. Sigma-square values for 11,000 iterations of the BART algorithm to predict log prices. Horizontal blue line represents the final average value of sigma-squared after the 10,000 burn-in iterations, and green lines show the 95 percent confidence interval.

⁴³For the technical aspects of this method, see Stekhoven and Bühlmann 2012.

the probability of being sold at each of the five venue categories, based on observable characteristics. Then, within each venue category, another random forest was calibrated to predict lot size.

To calibrate the price prediction algorithm, we randomly split the data in half for training and testing, then used BART to predict price based on observable characteristics, venue, and lot size. The BART procedure slightly improves on random forests by using an iterative Bayesian backfitting Markov chain Monte Carlo algorithm to fit residuals.⁴⁴ Implementing the `bartMachine` package in R, we chose to burn in the BART process for 10,000 iterations and develop the predictive model in 1,000 additional iterations.⁴⁵ Figure 6 justifies these choices by showing that the BART errors converged to a stationary distribution after 10,000 trials. Overall, the BART procedure yielded out-of-sample pseudo-R-square values of 89.3 percent and fits 96.3 percent of the test data within 95 percent confidence bands (Figure 7).

Both random forests and BART tend to underestimate the upper tail of the price distribution and overestimate the lower tail. This is to be expected since our




FIGURE 7. Out-of-sample BART prediction intervals: prediction intervals for log prices of out-of-sample test data, using the calibrated BART model. Blue dots indicate objects that were fit within a 95 percent confidence interval; red Xs indicate objects that were not. Axes are on log scales.

⁴⁴On backfitting, see Hastie and Tibshirani 2000.

⁴⁵For an in-depth discussion of this particular BART technique, see Kapelner and Bleich 2016.




FIGURE 8. BART residuals: residual plot for log-price predictions from calibrated BART model.

descriptive variables cannot capture the idiosyncratic qualities that would set them apart, especially for high-value items. Similarly, the lowest-quality objects may have some sort of unique damage or may be otherwise unattractive or uncollectible in ways we cannot systematically record. Nevertheless, Figure 8 shows that even with these caveats the residuals are well behaved.

Step 3: Imputation

For excavated items, sales venue and lot size must be predicted based on other observable characteristics. To account for the full distribution of possible sales outcomes, we ran Monte Carlo simulation of object sales separately for Dura Europos and Tell Bi'a. First, a venue/lot size combination was drawn for each excavated item using predicted probabilities from the random forest models calibrated in Step 2. Next, a final sale price was drawn from the posterior distribution of the conditional expected log price, using the BART model calibrated and tested on market data. The final output was the aggregate price of all objects; this procedure was iterated until the average aggregate price and the limits of the resulting 95 percent bootstrapped confidence interval all changed by less than 0.01 percent.

BIBLIOGRAPHY

Ashenfelter, Orley, and Kathryn Graddy. 2006. "Art Auctions." In *The Handbook of the Economics of Art and Culture*, edited by Victor Ginsburgh and David Throsby, vol. 1, 909–45. New York: Elsevier.

- Baird, Jennifer. 2012. "Dura Deserta: The Death and Afterlife of Dura-Europos." *Vrbes Extinctae: Archaeologies of Abandoned Classical Towns*, edited by Neil Christie and Andrea Augenti, 307–29. Burlington, VT: Ashgate.
- Baird, Jennifer. 2014. *The Inner Lives of Ancient Houses: An Archaeology of Dura-Europos*. Oxford: Oxford University Press.
- Beggs, Alan, and Kathryn Graddy. 2009. "Anchoring Effects: Evidence from Art Auctions." *American Economic Review* 99: 1027–39.
- Breiman, Leo. 2001. "Random Forests." *Machine Learning* 45: 5–32.
- Brodie, Neil. 1998. "Pity the Poor Middleman." *Culture without Context* 3: 7–9.
- Brodie, Neil. 2012. "Uncovering the Antiquities Market." In *The Oxford Handbook of Public Archaeology*, edited by Robin Skeates, Carol McDavid, and John Carman, 230–52. Oxford: Oxford University Press.
- Brodie, Neil, and Colin Renfrew. 2005. "Looting and the World's Archaeological Heritage: The Inadequate Response." *Annual Review of Anthropology* 34: 343–61.
- Brody, Lisa. 2011. "Yale University and Dura-Europos: From Excavation to Exhibition." In *Dura Europos: Crossroads of Antiquity*, edited by Lisa Brody and Gail Hoffman, 17–32. Chestnut Hill, MA: McMullen Museum of Art.
- Campbell, Peter. 2013. "The Illicit Antiquities Trade as a Transnational Criminal Network: Characterizing and Anticipating Trafficking of Cultural Heritage." *International Journal of Cultural Policy* 20: 131–53.
- Casana, Jesse. 2015a. "Satellite Imagery-based Analysis of Archaeological Looting in Syria." *Near Eastern Archaeology* 78: 142–52.
- Casana, Jesse, ed. 2015b. "Special Issue: The Cultural Heritage Crisis in the Middle East." *Near Eastern Archaeology* 78: 3.
- Casana, Jesse, and Elise Laugier. 2017. "Satellite Imagery-based Monitoring of Archaeological Site Damage in the Syrian Civil War." *PLOS One* 12, no. 11: e0188589.
- Casana, Jesse, and Panahipor Mitra. 2014. "Satellite-based Monitoring of Looting and Damage to Archaeological Sites in Syria." *Journal of Eastern Mediterranean Archaeology and Heritage Studies* 2: 128–51.
- Chanel, Olivier, Louis-André Gérard-Varet, and Victor Ginsburgh. 1996. "The Relevance of Hedonic Price Indices." *Journal of Cultural Economics* 20: 1–24.
- Chipman, Hugh, Edward I. George, and Robert E. McCulloch. 2010. "BART: Bayesian Additive Regression Trees." *Annals of Applied Statistics* 4: 266–98.
- Chippindale, Christopher, and David Gill. 2000. "Material Consequences of Contemporary Classical Collecting." *American Journal of Archaeology* 104: 463–511.
- Coggins, Clemency. 1972. "Archaeology and the Art Market." *Science* 175: 263–66.
- Contreras, Daniel A., and Neil Brodie. 2010. "Quantifying Destruction: An Evaluation of the Utility of Publicly-Available Satellite Imagery for Investigating Looting of Archaeological Sites in Jordan." *Journal of Field Archaeology* 35: 98–111.

- Cumont, Franz Valery Marie. 1925. *Fouilles de Doura-Europos (1922-1923)*. Paris: P. Geuthner.
- Danti, Michael. 2014. "Heritage in Peril: Syria and Iraq." Paper presented to the Metropolitan Museum of Art, New York, 22 September.
- Elia, Ricardo. 2001. "Analysis of the Looting, Selling, and Collecting of Apulian Red-Figure Vases: A Quantitative Approach." In *Trade in Illicit Antiquities: The Destruction of the Archaeological Heritage*, edited by Neil Brodie, Jennifer Doole, and Colin Renfrew, 145–53. Cambridge, UK: McDonald Institute for Archaeological Research.
- Elkins, Nathan. 2012. "The Trade in Fresh Supplies of Ancient Coins: Scale, Organization, and Politics." In *All the King's Horses: Essays on the Impact of Looting and the Illicit Antiquities Trade on Our Knowledge of the Past*, edited by Paula Lazarus and Alex Barker, 91–107. Washington, DC: Society for American Archaeology.
- Fanusie, Yaya, and Alexander Joffe. 2015 *Monumental Fight: Countering the Islamic State's Antiquities Trafficking*. Washington, DC: Foundation for the Defense of Democracies.
- Farchakh-Bajjal, Joanne. 2008. "Who Are the Looters at Archaeological Sites in Iraq?" In *Antiquities under Siege: Cultural Heritage Protection after the Iraq War*, edited by Lawrence Rothfield, pp. 49–56. New York: Altamira.
- Fay, Emily. 2011. "Virtual Artifacts: eBay, Antiquities, and Authenticity." *Journal of Contemporary Criminal Justice* 27: 449–64.
- Fearon, James. 2004. "Why Do Some Civil Wars Last Much Longer Than Others?" *Journal of Peace Research* 41: 275–301.
- Gibson, McGuire. 2003. "From the Prevention Measures to the Fact-Finding Mission." *Museum International* 55: 108–18.
- Haber, Stephen, and Victor Menaldo. 2011. "Do Natural Resources Fuel Authoritarianism? A Reappraisal of the Resource Curse." *American Political Science Review* 105: 1–26.
- Hastie, Trevor, and Robert Tibshirani. 2000. "Bayesian Backfitting (with Comments and a Rejoinder by the Authors)." *Statistical Science* 15: 196–223.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, Prediction*. New York: Springer.
- Ilan, David, Uzi Dahari, and Gideon Avni. 1989. "The Rampant Rape of Israel's Archaeological Sites." *Biblical Archaeological Review* 15: 38–41.
- James, Simon. 2014. In *Encyclopedia of Global Archaeology*, edited by Claire Smith, 2208–15. London: Springer.
- Kapelner, Adam, and Justin Bleich. 2016. "BartMachine: Machine Learning with Bayesian Additive Regression Trees." *Journal of Statistical Software* 70: 1–40. doi:10.18637/jss.v070.i04.
- Kersel, Morag. 2005. "Archaeology's Well Kept Secret: The Managed Antiquities Market." In *Symposium on Mediterranean Archaeology*, edited by Camilla Briault, Jack Green, Anthi Kaldelis, and Anna Stellatou, 79–83. Oxford: British Archaeological Reports.
- Kersel, Morag. 2006. "From the Ground to the Buyer: A Market Analysis of the Trade in Illegal Antiquities." In *Archaeology, Cultural Heritage and the Antiquities Trade*, edited by Neil Brodie,

- Morag Kersel, Christina Luke, and Kathryn Walker Tubb, 188–205. Gainesville: University Press of Florida.
- Kraft, Robert. 2007. "Pursuing Papyri and Papyrology by Way of eBay: A Preliminary Report." Paper presented at the Twenty-fifth International Congress of Papyrology, University of Michigan, Ann Arbor, 3 August.
- MacKenzie, Simon. 2009. "Identifying and Preventing Opportunities for Organized Crime in the International Antiquities Market." In *Organised Crime in Art and Antiquities*, edited by Stefano Manacorda, 41–62. Milan: International Scientific and Professional Advisory Council.
- MacKenzie, Simon, and Penny Green, eds. 2009. *Criminology and Archaeology: Studies in Looted Antiquities*. Oxford: Hart.
- Marrone, James V. 2018. "Quantifying the Supply Chain for Near Eastern Antiquities in Times of War and Conflict." *Journal of Cultural Heritage* (forthcoming).
- Massey, Douglas S., and Audrey Singer. 1995. "New Estimates of Undocumented Mexican Migration and the Probability of Apprehension." *Demography* 32: 203–13.
- McAndrew, Clare, James Smith, and Rex Thompson. 2012. "The Impact of Reserve Prices on the Perceived Bias of Expert Appraisals of Fine Art." *Journal of Applied Econometrics* 27: 235–52.
- Nørskov, Vinnie. 2002. "Greek Vases for Sale: Some Statistical Evidence." In *Illicit Antiquities: The Theft of Culture and the Extinction of Archaeology*, edited by Neil Brodie and Kathryn Walker Tubb, 23–37. London: Routledge.
- Parcak, Sarah, David Gathings, Chase Childs, Greg Mumford, and Eric Cline. 2016. "Satellite Evidence of Archaeological Site Looting in Egypt: 2002–2013." *Antiquity* 90: 188–205.
- Rose, Jerome, and Dolores L. Burke. 2014. "Making Money from Buried Treasure." *Culture without Context* 14: 4–8.
- Ross, Michael. 2006. "A Closer Look at Oil, Diamonds, and Civil War." *Annual Review of Political Science* 9: 265–300.
- Stekhoven, Daniel J., and Peter Bühlmann. 2012. "MissForest: Non-Parametric Missing Value Imputation for Mixed-Type Data." *Bioinformatics* 28: 112–18.
- Stone, Elizabeth. 2008. "Patterns of Looting in Southern Iraq." *Antiquity* 82: 125–38.
- Toll, Nicholas Peter. 1946. *The Excavations at Dura Europos: Preliminary Report of the Ninth Season of Work, 1935-1936*. Part II. The Necropolis. Yale: Yale University Press.
- Topçuoğlu, Oya and Tasha Vorderstrasse. Forthcoming. "Small Finds Add Up: Cylinder Seals and Coins from Iraq and Syria on the Online Market." *International Journal of Cultural Property*.
- United Nations Office on Drugs and Crime. 2005. "Estimating the Value of Illicit Drug Markets." In *World Drug Report*, vol. 1: Analysis, 123–46. Vienna: United Nations Office on Drugs and Crime.
- Vitelli, Karen D. 1979. "What Have You Done about the Antiquities Market Today?" *Journal of Field Archaeology* 6: 75–77.