

THE QUARTERLY JOURNAL OF ECONOMICS

Vol. 138

2023

Issue 3

AI-TOCRACY*

MARTIN BERAJA
ANDREW KAO
DAVID Y. YANG
NOAM YUCHTMAN

Recent scholarship has suggested that artificial intelligence (AI) technology and autocratic regimes may be mutually reinforcing. We test for a mutually reinforcing relationship in the context of facial-recognition AI in China. To do so, we gather comprehensive data on AI firms and government procurement contracts, as well as on social unrest across China since the early 2010s. We first show that autocrats benefit from AI: local unrest leads to greater government procurement of facial-recognition AI as a new technology of political control, and increased AI procurement indeed suppresses subsequent unrest. We show that AI innovation benefits from autocrats' suppression of unrest: the contracted AI firms innovate more both for the government and commercial markets and are more likely to export their products; noncontracted AI firms do not experience detectable negative spillovers. Taken together, these results suggest the possibility of sustained AI innovation under the Chinese regime: AI innovation entrenches the regime, and the regime's investment in AI for political control stimulates further frontier innovation. *JEL Codes:* O30, P00, E00, L5, L63, O25, O40.

*Many appreciated suggestions, critiques, and encouragement were provided by Tim Besley, Filipe Campante, Sergei Guriev, Peter Lorentzen, Torsten Persson, Nancy Qian, Imran Rasul, Andrei Shleifer, Jon Weigel, and many seminar and conference participants. We thank Yangsai Chen, Haoran Gao, Joelle Liu, Junxi Liu, Lingfei Liu, Xinze Liu, Shuhao Lu, Wenwei Peng, Wenbo Teng, Evelyn Xinmei Yang, Peilin Yang, and Tanggang Yuan for their excellent work as research assistants. Kao acknowledges financial support from the National Science Foundation Graduate Research Fellowship Program. Yang acknowledges financial support from the Harvard Data Science Initiative; Yuchtman acknowledges financial support from the British Academy under the Global Professorships program.

© The Author(s) 2023. Published by Oxford University Press on behalf of the President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2023), 1349–1402. <https://doi.org/10.1093/qje/qjad012>. Advance Access publication on March 13, 2023.

I. INTRODUCTION

Autocratic institutions have long been viewed as fundamentally misaligned with frontier innovation: autocrats' political and economic rents are eroded by technological change and economic growth; and incentives to innovate are stifled by threats and acts of expropriation under autocracy.

Recent scholarship, however, has suggested that artificial intelligence (AI) technology—considered to be the basis for a “fourth industrial revolution” (Schwab 2017)—may exhibit characteristics that allow an alignment between frontier innovation and autocracy. As a technology of prediction (Agrawal, Gans, and Goldfarb 2019), AI may be particularly effective at enhancing autocrats' social and political control (Zuboff 2019; Acemoglu 2021; Tirole 2021).¹ Furthermore, government purchases of AI may generate broad innovation spillovers, such as those observed among dual-use technologies (Moretti, Steinwender, and Van Reenen forthcoming). More specific to AI, because government data are inputs into developing AI prediction algorithms and can be shared across multiple purposes (Beraja, Yang, and Yuchtman forthcoming), autocracies' collection and processing of data for purposes of political control may directly stimulate AI innovation for the commercial market, far beyond government applications.² These arguments imply the possibility of a mutually reinforcing relationship in which governments procure AI to achieve political control, and this procurement stimulates further innovation in the technology.

Empirical evidence supporting such a mutually reinforcing relationship is lacking. As a technology still in its infancy, there exists little systematic evidence on the political deployment of

1. Predictions are extraordinarily valuable for an autocrat trying to maintain social and political control. They can serve to enhance monitoring (e.g., using prediction algorithms to identify and track individuals), to project human behaviors (e.g., identifying people who are more likely to engage in political unrest), and to shape behaviors (e.g., providing targeted sticks and carrots).

2. Moreover, government procurement may increase private data collection, which can then be shared across firms due to its nonrivalry (Aghion, Jones, and Jones 2019; Jones and Tonetti 2020). Procuring AI technologies could also stimulate innovation through traditional “crowding-in” channels, including the production of nontangible assets (e.g., ideas) and technological spillovers across government and commercial applications, within a firm and between firms. Public procurement also provides resources to firms that may allow them to cover fixed costs of innovation and overcome financial constraints.

AI,³ and essentially none on its efficacy in maintaining political control.⁴ Moreover, while [Beraja, Yang, and Yuchtman \(forthcoming\)](#) show that government data accessed through procurement contracts is valuable to stimulate commercial innovation among the contract-awarded firms, this does not imply that aggregate frontier innovation arises from contracts issued explicitly out of political-control motives.⁵

In the context of facial-recognition AI in China, we present evidence that frontier innovation and an autocratic regime can indeed be mutually reinforcing. In addition to the economic importance and geopolitical stakes of this context, it is also particularly suitable for studying innovation under autocracy. Maintaining political control is a paramount objective of the ruling Chinese Communist Party (see, among others, [Shirk 2007](#)). All citizens, even China's most successful entrepreneurs, are threatened by an unconstrained autocrat's ability to violate their property rights—and at times civil rights.⁶ Moreover, China is among the world's leading producers of commercial AI innovation, and facial recognition is one of the most important fields of AI technology.⁷

To conduct our empirical analyses, we combine several data sources: (i) episodes of local political unrest in China from the GDELT project; (ii) local public security agencies' procurement of facial-recognition AI (and complementary surveillance technology) primarily from China's Ministry of Finance; and (iii) China's facial-recognition AI firms' new software development (registered

3. Indeed, government use of AI is aimed at achieving a variety of objectives, from crime reduction to traffic control to environmental monitoring; see, among others, [Chin and Lin \(2022\)](#).

4. [Greitens \(2020\)](#) highlights that one of the most important open questions about surveillance AI is whether its deployment actually enhances political repression.

5. For example, AI firms selected to provide government services in politically sensitive environments may be differentially specialized in technology (only) for government use; firms' production and service provision for government in such a context may require significant reallocation of resources that could crowd out their commercial and broader innovation activities; contract-awarded firms may impose substantial negative spillovers on peer firms that have not received such contracts due to business stealing or attracting productive inputs such as human capital.

6. For example, Jack Ma, the founder of Alibaba, was detained for months upon arousing the ire of the Chinese Communist Party. For example, see [Yang and Wei \(2020\)](#).

7. For example, in 2020, computer vision was the second largest field of study in AI by publications on arXiv, accounting for 31.7% of the total publications ([Zhang et al. 2021](#)).

with the Ministry of Industry and Information Technology), as well as their software export deals (compiled through press releases, news reports, and other sources). Linking data sets (i) and (ii) allows us to test whether autocracies procure facial-recognition AI for purposes of political control, whether facial-recognition AI is effective in suppressing unrest, and whether AI procurement is associated with complementary changes in the technology of political control (such as the procurement of surveillance cameras). Then, linking these two data sets to (iii) enables us to test the extent to which facial-recognition AI innovation benefits from politically motivated procurement—specifically, whether total software development increases, whether software intended for commercial markets (beyond political uses) increases, and whether internationally competitive products emerge.⁸

We begin by examining the first direction in a mutually reinforcing relationship: whether AI technology can effectively enhance autocrats' political control. We test whether autocrats respond to political unrest by procuring facial-recognition AI technology. We find that indeed they do: locations experiencing episodes of political unrest increase their public security procurement of facial-recognition AI. This result holds controlling for a range of time-varying local characteristics, including local government fiscal revenue. One might wonder whether procuring public security AI was already on a different trend in locations experiencing political unrest (e.g., due to anticipation of subsequent unrest or because of different rates of economic growth). However, we find quantitatively small increases in AI procurement just before episodes of political unrest, and a substantially larger increase in AI procurement during the quarter immediately following episodes of unrest. One might also wonder whether time- and space-varying shocks are correlated with the occurrence of political unrest and with public security AI procurement. To address this concern, we implement an IV strategy exploiting variation in the occurrence of political unrest arising from local weather conditions, and we find qualitatively and quantitatively similar results. Further evidence of a broad technological upgrade in

8. This may be more precisely called “product innovation at the technological frontier” than “frontier innovation,” given that new software products do not necessarily imply that a firm is substantially shifting the frontier. For brevity, we refer to novel production of a frontier technology as “frontier innovation,” but we acknowledge that these new products we study are generally examples of micro-inventions rather than macro-inventions, to use the terminology of Mokyr (1990).

political control suggests that the increased AI procurement reflects an active choice by public security agencies in response to political unrest, rather than mere “window dressing.” We find that locations experiencing political unrest purchase more high-resolution surveillance cameras, which provide the crucial data input for facial-recognition technology. Moreover, public security agencies that have procured more facial-recognition AI technologies not only reduce their subsequent hiring of police staff but also shift the composition of the police force toward higher-skilled desk jobs that complement the use of AI technology.

Local governments’ purchases of AI technology in response to the occurrence of political unrest suggest at least a belief in the effectiveness of such technology in curbing future unrest. We study whether the increased AI procurement enhances autocrats’ political control. Precisely because AI is procured endogenously in locations susceptible to political unrest, instead of examining the relationship between AI procurement and subsequent local unrest, we examine how past investment in public security AI mitigates the effect of exogenous shocks that tend to instigate political unrest. We find that weather conditions conducive to unrest have smaller effects on contemporaneous unrest in prefectures that have accumulated a larger stock of public security AI capacity up to the previous quarter. Conducting a placebo exercise, we find that such a relationship is not observed in response to the accumulated non-public security AI capacity, suggesting that our results are driven by the deployment of public security AI per se, rather than by differing socioeconomic conditions in politically sensitive contexts. Importantly, our results are not due to the time-varying effects of past unrest that are associated with public security AI investment: local experience of past unrest is not associated with differential unrest arising from current weather conditions.

Having established that AI does strengthen autocrats’ political control, we examine the second direction in a mutually reinforcing relationship: whether politically motivated AI procurement stimulates AI innovation. We define politically motivated AI procurement as purchases by public security agencies in prefectures that experienced above-median levels of political unrest in the previous quarter.⁹ We estimate an event study

9. Public security contracts issued following episodes of unrest are proxies for contracts with an underlying political motivation. We acknowledge that such motivation is not explicitly observed. For brevity, we refer to these contracts as politically motivated.

specification, estimating the effects of a firm's first politically motivated contract on AI software production, controlling for firm and time period fixed effects. We find that prior to receipt of a politically motivated contract, firms do not exhibit differential software production, suggesting that selection into such contracts is largely accounted for by the firm fixed effects. Within the first year of contract receipt, firms produce significantly more AI software; by two years postcontract, they produce around 10 (48.6%) more software products. Such an increase is observed not just among software intended for government uses, but importantly among software intended for broader commercial applications. To address the concern that political unrest is more likely to occur in economically dynamic locations where commercial AI innovation is also greater, we instead identify politically sensitive environments and classify politically motivated procurement contracts using predicted political unrest based on weather conditions, and the results are qualitatively unchanged. In other words, plausibly exogenous episodes of political unrest promote commercial AI innovation through increased local public security demand for AI.

An important concern is whether contracts issued in a politically sensitive environment either provide benefits to firms that could account for their increase in innovation activities (e.g., closer government–business ties) or induce differential selection of contract-awarded firms (e.g., greater scrutiny of firms' capacity and potential). To address this concern, we compare the effects of public security contracts issued in this politically sensitive environment—these contracts are most plausibly politically motivated—with the effects of non–public security contracts issued in the same environment. This allows us to isolate the effects of politically motivated contracts beyond the consequences arising from generic contracts issued in a politically sensitive environment. Using a triple-differences empirical strategy, we find that receipt of a politically motivated public security contract is associated with significantly greater innovation of commercial (as well as government) software, relative to the receipt of a contract with a non–public security arm of the government. We find no evidence of differential precontract trends in software innovation, supporting a causal interpretation of our findings. To establish the international competitiveness of the new AI software produced following politically motivated contracts, we test whether receipt of such contracts is associated with a greater likelihood that firms export their products. Indeed, we find a tripling of the likelihood that firms begin exporting, suggesting that politically motivated

contracts have pushed the contracted firms to the technological frontier.

Finally, we investigate whether autocrats' politically motivated AI demands distort the trajectory of innovation at the firm level and at the aggregate level. We ask whether, at the firm level, politically motivated public security contracts induce less commercial innovation than similar public security contracts issued in politically neutral environments (namely, when not preceded by local unrest occurrence). We find that the effects of politically motivated public security contracts on commercial AI innovation are not smaller than other, politically neutral contracts (if anything, we find the effects are larger), suggesting that the political motivation does not diminish the effect of procurement contracts on AI firms' commercial innovation. We consider the possibility that our firm-level findings may be offset in the aggregate by negative spillovers to other firms, for example, due to the allocation of resources or business-stealing effects. Specifically, we examine AI innovation among firms never receiving procurement contracts that are: (i) headquartered in localities that have experienced political unrest; (ii) headquartered in localities where AI firms receiving politically motivated contracts are also headquartered; or (iii) part of a mother firm with other subsidiaries that have received politically motivated contracts. We find no evidence of negative spillovers in these cases. If anything, we observe positive spillovers to firms not receiving contracts. This suggests that our firm-level effects of politically motivated contracts may increase frontier AI innovation in the aggregate, and such an aggregate effect may be larger than the total firm-level effects that we identify among contracted firms.

Taken together, these results imply that China's autocratic political regime and the rapid innovation in its AI sector are not in conflict but mutually reinforce each other.¹⁰ When such a mutually reinforcing relationship is sufficiently strong to overcome distortions in autocracies that discourage innovation (e.g., risk of expropriation), it could support an equilibrium—"AI-tocracy"—where an autocratic regime is entrenched, and frontier AI

10. We do not interpret our findings as indicating that China's political stability is primarily achieved through AI technology (yet), nor that China's AI innovation is primarily rooted in political repression. Rather, our findings suggest that a component of China's coercive capacity is derived from the application of AI technology, and China's political repression in turn contributes to AI innovation and in part leads to the rise of China as a leading innovator in AI.

innovation is sustained.¹¹ It does so by generating a perpetuating cycle in which autocrats are strengthened by AI innovation, and their procurement of this innovation stimulates further innovation, which in turn further strengthens the autocrats.

More generally, our analysis of the forces that support a mutually reinforcing relationship between autocrats' political repression and frontier innovation in facial-recognition AI sheds light on other prominent episodes of frontier innovation under nondemocratic regimes. Such episodes—ranging from the development of aerospace technology in the Soviet Union to chemical engineering innovation in imperial Germany—are difficult to reconcile with the large literature that highlights forces that limit innovation and growth in nondemocratic contexts.¹² These episodes, however, share important features that mirror China's facial recognition AI sector: first, the nondemocratic regimes appear to derive political power from frontier innovation; second, recognizing the political benefits of innovation, the regimes provide financial and institutional support that may be instrumental to technological development. To the extent that these mutually reinforcing forces overcome traditional autocratic frictions, innovation can entrench autocracies and be promoted by them in a sustained manner.¹³

Our work relates to several additional areas of very active research. We contribute to a growing literature on the socioeconomic consequences of AI technology. Over the past several years, many economists have been studying the far-reaching consequences of an emerging AI-led economy. However, much of

11. In fact, autocrats' sustained demand for AI technology for political control may also enhance their ability to commit to protect AI innovators' property rights, thus reducing the risk of expropriation.

12. The effects of political institutions on economic growth and frontier innovation have been studied by, among others, North and Weingast (1989), Acemoglu and Robinson (2006, 2012), Aghion, Alesina, and Trebbi (2008), and North, Wallis, and Weingast (2009). Autocracies may also exhibit reduced innovation due to corruption and the misallocation of talent (Murphy, Shleifer, and Vishny 1989; Shleifer and Vishny 2002). The effects of economic growth on political institutions have been studied by Lipset (1959), Barro (1996), and Glaeser, Ponzetto, and Shleifer (2007) (see Treisman 2020 for a review).

13. One also observes examples of mutually reinforcing relationships between democratic regimes and frontier innovation. One prominent case is the military innovation developed by DARPA in the United States, and its well-known commercial innovation consequences (e.g., the internet). We do not argue that innovation only supports autocratic regimes but that such a regime-enhancing effect of technology may be particularly relevant in nondemocracies due to their otherwise unfavorable environment for innovation.

the literature focuses on the economic consequences of AI: from its impact on the labor market (Acemoglu and Restrepo 2019, 2019) and how governments should respond to it (Beraja and Zorzi 2022), to how it reshapes market power and competition (Jones and Tonetti 2020; Eeckhout and Veldkamp 2022), to how it changes global trade (Goldfarb and Treffer 2018), to how it affects socioeconomic inequality (Korinek and Stiglitz 2018) and economic growth (Aghion, Jones, and Jones 2019; Farboodi and Veldkamp 2022). Some recent research has considered the social consequences of AI, in particular, discrimination arising from the potential biases in its algorithms (Kleinberg et al. 2018; Cowgill and Tucker 2020). This article provides the first direct evidence on the political consequences of AI technology: it can produce more effective political control, potentially entrenching autocratic government.

We also contribute to the literature on the role of state capacity in economic development (e.g., Besley and Persson 2009). The mutually reinforcing relationship we observe between a regime and frontier innovation can also be observed in settings beyond autocracies where the state exercises its fiscal capacity to support frontier technology (e.g., DARPA in the United States). We highlight the possibility of sustained innovation arising from an autocrat's exertion of state capacity for political control. Thus, we contribute to the literature allowing for the possibility of growth under extractive institutions (e.g., Acemoglu and Robinson 2020; Dell and Olken 2020).¹⁴ Beraja, Yang, and Yuchtman (forthcoming) find that Chinese government contracts stimulate AI innovation, but do not determine whether such contracts strengthen the autocrats, and whether politically motivated contracts in particular can foster commercial innovation. In this article, we demonstrate that frontier innovation can be sustained in autocracy as a result of their mutually reinforcing relationship. In fact, this implies a political economy trajectory that defies conventional wisdom: the Chinese case suggests a stable equilibrium

14. In addition to works cited here, a large empirical literature identifies negative effects of extractive institutions on long-run development (e.g., Acemoglu, Johnson, and Robinson 2002; Nunn 2008; Dell 2010; Lowes and Montero 2021). There has been a small strand of the literature that documents the positive economic consequences of colonial investments, particularly in transportation infrastructure and human capital (e.g., Huillery 2009; Cagé and Rueda 2016; Donaldson 2018; Valencia Caicedo 2019).

exhibiting sustained frontier innovation and further entrenched autocracy.¹⁵

We thus add to a large literature on the relationship between technology and political stability. Recent papers find that advances in information and communication technologies and the diffusion of social media have supported protest movements and populist parties in a broad range of settings (Campante, Durante, and Sobrio 2018; Enikolopov, Makarin, and Petrova 2020; Qin, Strömberg, and Wu 2020; Guriev, Melnikov, and Zhuravskaya 2021). We, on the other hand, contribute to a literature that documents how technological change can repress political unrest, thus strengthening autocracies and incumbents more generally. This literature describes the evolution of repressive technology: Autocracy 1.0—the state as a monopolist of violence using the threat of brute force to produce compliance out of fear (Olson 1993); Autocracy 2.0—the state as manipulator of information using propaganda and censorship to produce compliance out of persuasion (Cantoni et al. 2017; Roberts 2018; Chen and Yang 2019; Guriev and Treisman 2019); and finally, Autocracy 3.0—the state (and its AI) as monitor, predictor, and manipulator of behaviors to produce compliance using targeted behavioral incentives (Tirole 2021). To this literature, we provide the first empirical evidence on the systematic deployment of AI as a part of the state's political control apparatus, documenting its procurement alongside complementary technological inputs and its effects on maintaining political stability.¹⁶

Finally, we contribute to the literature on the political economy of growth in China. While much work emphasizes factors that promote China's growth despite its autocratic politics (Lau, Qian, and Roland 2000; Brandt and Rawski 2008; Song, Storesletten, and Zilibotti 2011), we join an emerging strand of the literature that highlights China's autocratic institutional features that facilitate growth (Bai, Hsieh, and Song 2020). Importantly, we demonstrate that China's stimulus of facial-recognition AI innovation is not due to marginal improvements in institutional dimensions

15. It is important to note that this political-economy equilibrium is not inevitable, because the mutually reinforcing relationship may be offset by autocratic distortions (e.g., risks of expropriation).

16. Our findings of AI technology being deployed in response to political unrest also contribute to a growing literature that studies authoritarian responsiveness to citizens' political grievances (e.g., Tsai 2007; Chen, Pan, and Xu 2016; Campante, Chor, and Li forthcoming).

such as protection of property rights and rule of law, nor to the enhancement of infrastructure or state capacity more generally. Instead, AI innovation is spurred directly by the application of political repression itself.

In [Section II](#), we describe the empirical context and the data sources we use. In [Section III](#), we present evidence of the effects of AI technology on autocratic political control. In [Section IV](#), we present the evidence on the effects of politically motivated procurement of AI on innovation. Finally, in [Section V](#), we conclude by discussing the implications of our findings.

II. EMPIRICAL CONTEXT AND DATA

AI technologies have been argued to have characteristics that could generate a mutually reinforcing relationship between innovation and autocracy.¹⁷ We test for the two directions of a mutually reinforcing relationship between autocracy and frontier innovation in the context of facial-recognition AI technology in China.

To test whether frontier AI innovation enhances autocratic political control (the first direction of the mutually reinforcing relationship), we examine (i) whether AI procurement is motivated by the regime's desire for political control, and (ii) whether procurement of AI technology out of political motivation indeed enhances the regime's political control by reducing unrest.

The Chinese regime is particularly concerned with political unrest ([Shirk 2007](#); [King, Pan, and Roberts 2013](#)). Thus, consider local government officials' response to an episode of local unrest. Anticipating that such unrest may persist into subsequent periods—either due to socioeconomic shocks that are serially correlated or because unrest participation itself is path dependent ([Madestam et al. 2013](#); [Bursztyn et al. 2021](#))—the local officials may procure facial-recognition AI technology and upgrade their political-control technology. Such technology could allow the government to preemptively identify, crack down on, and deter the participants in future unrest, thus mitigating the effect of future shocks on the occurrence of local unrest.

17. While we focus on the AI sector in this article, mutually reinforcing relationships between autocracy and frontier innovation appear to have been present in other prominent historical episodes. In [Online Appendix A](#), we describe several such episodes, including the success of scientific innovation in the Soviet Union and the emergence of the second German empire as a powerhouse of science, industry, and innovation.

To test the second direction of the mutually reinforcing relationship, we examine whether politically motivated procurement of AI technology stimulates frontier AI innovation by the firms awarded contracts. Government procurement could provide these firms with valuable inputs, such as access to rich public security data and revenue streams, which may allow the AI firms to develop more and newer AI products. To the extent that these inputs may be shared across multiple purposes, the AI firms could increase their innovation activities in the commercial sector above and beyond products developed for government purposes. We close by considering the effects of politically motivated contracts on firms not awarded these contracts, thus gauging the potential aggregate innovation consequences.

To conduct our empirical analyses, we combine data on (i) episodes of local political unrest in China; (ii) local governments' procurement of facial-recognition AI technology and complementary technology for political control; and (iii) facial-recognition AI firms' software innovation and product export activities. We also describe auxiliary data sources used for various empirical exercises in [Online Appendix B](#).

II.A. Political Unrest

We collect data on political unrest from the Global Database of Events, Language, and Tone (GDELT) Project. The GDELT project records instances of events based on articles from a comprehensive, global set of news feeds.¹⁸ We restrict our analysis to events taking place in China between 2014 and 2020.¹⁹ In sum, we find 9,267 events indicating political unrest, corresponding to three broad categories: protests, demands, and threats.²⁰ [Figure I, Panel A](#), presents the spatial distribution of the political unrest

18. Text analysis and machine-learning methods are applied to the contents of these articles to identify salient characteristics, such as event location (which we geocode at the prefecture level), date of the event, and the nature of these events. See <https://www.gdeltproject.org> for a detailed description of the GDELT Project and its methodology.

19. The GDELT Project greatly expanded their scope of sources and text analysis capabilities in 2014, making coverage before 2014 less complete and reliable. From 2014 to 2020, there are over 100 news sources that provide coverage on China. When multiple news sources cover the same event, GDELT records only one event.

20. Each event is classified under the Conflict and Mediation Events Observations (CAMEO) event and actor codebook, in which protests (e.g., demonstrations, hunger strikes for leadership change), demands (e.g., demands for material aid,

Panel A: Unrest events



Panel B: Public security AI contracts

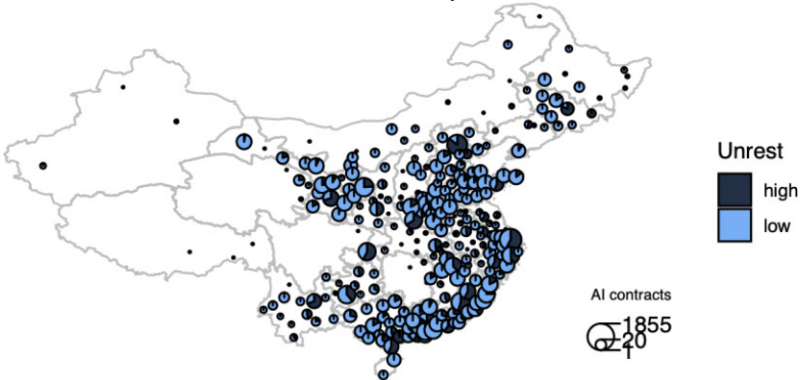


FIGURE I

Unrest and Public Security AI Contracts

Each circle represents a prefecture in our data set that has at least one public security AI contract that is an AI firm's first government contract. In Panel A, circle size indicates the number of unrest events in a prefecture; in Panel B, circle size indicates the number of public security AI contracts awarded in the prefecture (larger circles indicate more, log scale). Circle shading in Panel B indicates the fraction of first AI contracts that were procured during high- or low-unrest periods, where the within-prefecture variation comes from changes in the number of unrest events in a prefecture over time (a larger fraction of dark shading indicates a larger fraction of prefecture contracts procured during high-unrest periods).

TABLE I
SUMMARY STATISTICS

	Mean (1)	Std. dev. (2)
Panel A: Political unrest		
All events (per prefecture-quarter)	2.419	18.490
Protests	0.607	4.603
Demands	0.720	5.009
Threats	1.092	9.479
Panel B: Procurement of AI and the technology of political control		
All AI contracts (per prefecture-quarter)	3.976	7.818
Non-public security contracts	2.285	5.118
Public security contracts	1.691	3.476
First public security contracts	0.082	0.327
Surveillance cameras (per prefecture-quarter)	2,118	12,684
Police hires (per prefecture-year)	59.278	84.991
Panel C: Innovation of AI firms (flow)		
All software (per firm-quarter)	5.756	7.124
Government software	1.724	3.337
Commercial software	2.353	3.675
Panel D: Innovation of AI firms (cumulative, precontract)		
All software (per firm)	22.105	33.004
Government software	6.266	11.738
Commercial software	9.333	15.936

Notes. This table presents summary statistics at the prefecture-quarter level (firm-quarter and firm level for Panels C and D) for variables of interest. Column (1) shows the sample mean and column (2) the standard deviation. Panel A presents counts of unrest events, Panel B presents counts of local government-procured facial-recognition AI contracts and other technologies of political control, Panel C presents counts of software produced by facial-recognition AI firms per quarter (a flow variable), and Panel D presents counts of cumulative software produced by facial-recognition-AI firms up to the quarter before earning a contract. All software is equal to government software + commercial software + general AI software. For Panels A and B, $N = 8,167$ (Panel B police hires, $N = 2,672$). For Panel C, $N = 23,697$. For Panel D, $N = 5,462$.

that occurred during 2014 to 2020 in prefectures with AI contracts that we study; [Table I](#), Panel A, presents basic summary statistics of these political unrest events.

Given the state control of Chinese media sources, it is important to consider the possible effect of censorship on the quality of the GDELT data. We believe that the GDELT data are well-suited for our purposes for several reasons. First, the local unrest

leadership change, or policy change), and threats (e.g., threats to boycott, political dissent) are 3 of 20 top-level “verbs” that an event can be classified under, with the latter being relatively less politically threatening. We exclude a small number of events that occur at a national or international level. We are able to cross-check the unrest data against similar event counts from alternative sources, such as Radio Free Asia ([Qin, Strömberg, and Wu 2020](#)), and find very similar levels.

that we focus on has generally not been targeted for censorship by the Chinese authorities (Qin, Strömberg, and Wu 2017); some have even argued that media reporting on local unrest is particularly helpful for resolving the information asymmetry between the central and local government (Lorentzen 2014). Moreover, the GDELT data include a range of unrest events that differ in their political sensitivity, allowing us to examine whether the patterns we observe vary by political sensitivity.

1. *Local Weather Conditions Used to Construct Instruments for Political Unrest.* We use historical weather data originally collected by the World Meteorological Organization (WMO) and hosted by the National Oceanic and Atmospheric Administration (NOAA). Data are reported at the weather station–day level. These weather stations provide a wide variety of data at the daily level, including mean temperature, amount of precipitation, presence of fog, rain, hail, thunder, maximum windspeed recorded, and visibility.²¹ Importantly, we use all 18 weather variables that are consistently available around the world during our sampling period. We assign data to prefectures using the closest weather station to the given prefecture. For the 344 prefectures in our data set, this results in 260 unique weather stations whose data we use.

II.B. Procurement of AI and the Technology of Political Control

To observe the Chinese government's demand for AI technology, we extract information on 2,997,105 procurement contracts issued by all levels of the Chinese government between 2013 and 2019 from the Chinese Government Procurement Database, maintained by China's Ministry of Finance.²² The database contains information on the good or service procured, the date of the contract, the monetary size of the contract, the winning bid, as well as, for a subset of the contracts, information on bids that did not win.

21. Other weather variables are mean dew point, mean sea level (air) pressure, mean station pressure, mean wind speed, maximum wind gust, maximum temperature, minimum temperature, snow depth, and presence of tornadoes or funnel clouds. This weather data ranges from 2012 to 2020. There are a small number of observations for which weather data are missing (less than 1% of the total). For these, we impute data from the geographically nearest weather station, or in the one case when all stations are missing data on a given day, we take data from the following day and the same station instead.

22. See [Online Appendix](#) Figure A.I for an example contract.

To narrow our focus on the subset of contracts that procure facial-recognition AI technology such as data-processing services or platform solutions, we match the contracts with a list of facial-recognition AI firms. We identify (close to) all active firms based in China producing facial-recognition AI using information from Tianyancha, a comprehensive database on Chinese firms licensed by China's central bank.²³ We extract firms that are categorized as facial-recognition AI producers by the database, and we validate the categorization by manually coding firms based on their descriptions and product lists. We collect an array of firm-level characteristics such as founding year, capitalization, major external financing sources, as well as subsidiary and mother firm information. Overall, we identify 7,837 Chinese facial-recognition AI firms.²⁴

Our empirical exercises in particular concern the AI procurement contracts awarded by public security agencies of the Chinese government.²⁵ As an example from our data set, consider a contract signed between an AI firm and a municipal police department in Heilongjiang Province to "increase the capacity of its identity information collection system" on August 29, 2018. The contract specifies that the AI firm shall provide a facial-recognition system that should cover at least 30 million individuals, suggesting the large scale of data collection and processing that is required. In total, we identify 28,023 public security-related procurement contracts issued to AI firms.²⁶ They include the following four types of public security contracts from the Chinese

23. A primary source of firms' information compiled by Tianyancha is the National Enterprise Credit Information Publicity System, maintained by China's State Administration for Industry and Commerce. See [Online Appendix](#) Figure A.II for an example entry. We complement the Tianyancha database with information from Pitchbook, a database owned by Morningstar on firms and private capital markets around the world. See [Online Appendix](#) Figure A.III for an example entry.

24. These firms fall into three categories: (i) firms specialized in facial-recognition AI (e.g., Yitu); (ii) hardware firms that devote substantial resources to develop AI software (e.g., Hik-Vision); and (iii) a small number of distinct AI units in large tech conglomerates (e.g., Baidu AI).

25. Parts of our empirical strategy compare public security procurement contracts of AI to those awarded by non-public security units in the public sector, such as (public) banks, hospitals, and schools. There are a total of 6,557 non-public security related procurement contracts awarded to AI firms.

26. We present the cumulative number of AI procurement contracts in [Online Appendix](#) Figure A.IV (top panel), as well as the flow of new contracts signed

Government Procurement Database: (i) all contracts for China's flagship surveillance/monitoring projects—Skynet Project, Peaceful City Project, and Bright Transparency Project; (ii) all contracts with local police departments; (iii) all contracts with the border control and national security units; and (iv) all contracts with the administrative units for domestic security and stability maintenance, the government's political and legal affairs commission, and various “smart city” and digital urban management units of the government. Importantly, each contract is linked to a specific prefectural government buyer, and for the baseline analysis, we exclude those signed with the central or provincial government. Many firms receive multiple public security contracts; overall, 1,095 (14%) of facial-recognition AI firms in our data set receive at least one contract.²⁷ Figure I, Panel B, presents the spatial distribution of the facial-recognition AI contracts issued by public security units of the prefectural government.²⁸

In addition to the public security agencies' procurement of AI technology, we also collect information on a key complementary technology for political control: high-resolution surveillance cameras procured by the same agencies. These cameras, once deployed in the public space, could provide richer data that would make the facial-recognition AI platform more effective, and may also deter civilian unrest. Table I, Panel B, presents basic summary statistics of the facial-recognition AI procurement contracts issued by public security and non-public security agencies, as well as the procurement of surveillance cameras.

II.C. AI Firms Innovation

1. *Product Innovation: AI Software Development.* We collect all software registration records for our facial-recognition AI firms from China's Ministry of Industry and Information Technology, with which Chinese firms are required to register new software releases and major upgrades. We are able to validate our measure

in each month (bottom panel). Both public security and non-public security AI contracts have steadily increased since 2013.

27. This means that conditional on receiving a contract, on average a firm receives 25.6 contracts over our sampling period.

28. Some public security AI contracts are issued at the provincial level: for example, almost 40% of the public security AI contracts in Xinjiang are issued by the provincial government. Online Appendix Figure A.V plots the spatial distribution of public security AI contracts issued by either provincial or prefectural governments.

of software releases (using a single large firm), by cross-checking our data against the IPO Prospectus of MegVii, the world's first facial-recognition AI company to file for an IPO.²⁹ We find that our records' coverage is comprehensive (at least in the case of MegVii): MegVii's IPO Prospectus contains 103 software releases, all of which are included in our data set.

The count of new software releases (and major upgrades) represents product innovation.³⁰ Reflecting the economic value of such innovation, we observe that facial-recognition AI firms that develop more software have significantly and substantially higher market capitalization (see [Online Appendix Figure A.VI](#)).

We use a recurrent neural network (RNN) model with tensorflow—a frontier method for analyzing text using machine learning—to categorize software products according to their intended customers and (independently) by their function. Our categorization by customer distinguishes between software products developed for the government (e.g., “smart city—real-time monitoring system on main traffic routes”) and software products developed for commercial applications (e.g., “visual recognition system for smart retail”). We allow for a residual category of general application software whose description does not clearly specify the intended user (e.g., “a synchronization method for multi-view cameras based on FPGA chips”). By coding as “commercial” only those products that are specifically linked to commercial applications, and excluding products with ambiguous use, we aim to be conservative in our measure of commercial software products.

Our categorization by function first identifies software products that are directly related to AI (e.g., “a method for pedestrian counting at crossroads based on multi-view cameras system in complicated situations”). In the category of AI software, we also separately identify a subcategory of software that involve components related to surveillance (e.g., “tool that allows parents to locate and track lost children”).

To implement the two dimensions of categorization using the RNN model, we manually label 13,000 software products to

29. Source: Hong Kong Stock Exchange, <https://go.aws/37GbAZG>.

30. The National Science Foundation defines product innovation as “the market introduction of a new or significantly improved good or service with respect to its capabilities, user-friendliness, components, or sub-systems” in its Business Enterprise Research and Development Survey (see <https://www.nsf.gov/statistics/srvyberd>). See also [Bloom et al. \(2020\)](#).

produce a training corpus. We use word-embedding to convert sentences in the software descriptions into vectors based on word frequencies, where we use words from the full data set as the dictionary. We use a long short-term memory (LSTM) algorithm, configured with two layers of 32 nodes. We use 90% of the data for algorithm training, and 10% is retained for validation. We run 10,000 training cycles for gradient descent on the accuracy loss function. The categorizations perform well in general: we are able to achieve 72% median accuracy in categorizing software customer and 98% median accuracy in categorizing software function. [Online Appendix Figure A.VII](#) shows the summary statistics of the categorization output by customers and by function; and, [Online Appendix Figure A.VIII](#) presents the confusion matrix (type I and type II errors) of the predictions relative to categorization done by humans.³¹ [Table I](#), Panel C, presents basic summary statistics of the software innovation of all AI firms (regardless of whether they have received procurement contracts), and Panel D presents the cumulative AI software production prior to firms' receipt of their first public security procurement contracts.

2. *Frontier Technology: Firms AI Software Exports.* We construct a database of global AI trade deals using the bibliography of the Carnegie Endowment for International Peace's report *The Global Expansion of AI Surveillance* ([Feldstein 2019](#)). This bibliography focuses on international procurement of AI surveillance technology by governments, containing 1,300 citations spanning 75 countries.³² Examples of such deals include "Safe City Service Brings the Future to Laos: Huawei case studies" (China exporting to Laos in 2015), "Bosch equips Hong Kong-Zhuhai-Macao Bridge with customized security solutions" (Germany exporting to China in 2018), and "Digital Intelligence Is Helping Brazil's Federal Police Seize Millions in Assets to Bring Down Drug-Smuggling Kingpins" (Israel exporting to Brazil in 2020).

31. [Online Appendix Table A.I](#) presents the top words (in terms of frequency) used for the categorization. [Online Appendix Figure A.IX](#) presents the density plots of the algorithm's category predictions. The algorithm is very accurate in categorizing software for government purposes. The algorithm is relatively conservative in categorizing software products for commercial customers, and relatively aggressive in categorizing them as general purpose. In setting our categorization threshold for commercial software, we again aim to be conservative in our measure of commercial software products.

32. The original bibliography is accessible at https://www.zotero.org/groups/2347403/global_ai_surveillance/library.

We match each trade deal to the Chinese AI firms in our data, allowing us to identify the date from which the firms begin to export their AI products. For each Chinese AI firm, we also search through their press releases and news reports covering them to expand our database of AI trade deals. Among the 7,837 Chinese facial-recognition AI firms we study, we identify a total of 176 export deals.

III. THE ROLE OF AI IN AUTOCRATS' POLITICAL CONTROL

III.A. *The Effect of Political Unrest on AI Procurement and the Technology of Political Control*

Our empirical analyses begin by examining whether AI technology can effectively entrench autocrats. Specifically, we first test whether local public security agencies (e.g., police forces) respond to episodes of local political unrest by procuring more facial-recognition AI in the following quarter. The time lag reflects the administrative procedure and time needed to initiate and issue a contract in response to an event. We estimate panel models that control for locality and time period fixed effects, using OLS and implementing an IV specification that exploits differences in unrest occurrence due to local weather conditions.

We first describe the panel OLS strategy, where we estimate the following model:

$$(1) \quad AI_{i,t-1} = \beta Unrest_{it} + \alpha_t + \gamma_i + \delta_t X_i + \epsilon_{it},$$

where the explanatory variable of interest is $Unrest_{it}$, the local political unrest in prefecture i in quarter t , and $AI_{i,t-1}$ is the public security facial-recognition AI procurement per capita of prefecture i in the subsequent quarter. We control for time period and prefecture fixed effects, as well as different combinations of time-varying effects of prefecture socioeconomic characteristics. Standard errors are clustered at the prefecture level.

We present the results in [Table II](#), Panel A. To account for changing local economic and political conditions that may be related to both unrest occurrence and facial-recognition AI procurement, we control for the prefecture GDP (measured yearly) interacted with a full set of (quarterly) time fixed effects (column (1)), the prefecture's log population interacted with a full set of time fixed effects (column (2)), the prefectural government's annual fiscal revenue interacted with a full set of time fixed effects

TABLE II
THE EFFECT OF UNREST EVENTS ON FACIAL-RECOGNITION AI PROCUREMENT

	Public security AI procurement				
	(1)	(2)	(3)	(4)	(5)
Panel A: OLS					
Unrest events _{<i>t</i>-1}	0.199*** (0.043)	0.196*** (0.046)	0.199*** (0.044)	0.204*** (0.046)	0.205*** (0.044)
Panel B: LASSO IV					
Unrest events _{<i>t</i>-1}	0.377*** (0.084)	0.377*** (0.084)	0.377*** (0.084)	0.349*** (0.080)	0.348*** (0.080)
GDP × quarter	Yes	No	No	No	Yes
Log population × quarter	No	Yes	No	No	Yes
Gov. revenue × quarter	No	No	Yes	No	Yes
AI stock _{<i>t</i>-2}	No	No	No	Yes	Yes

Notes. This table presents regressions at the prefecture-quarter level. The outcome is the number of public security facial-recognition AI contracts procured by the local government, standardized to mean = 0 and variance = 1. The explanatory variable of interest is the occurrence of unrest events in the corresponding prefecture during the preceding quarter, also standardized. Column (1) controls for prefecture GDP × quarter effects, column (2) controls for log prefecture population × quarter effects, column (3) controls for prefectural government tax revenue × quarter effects, column (4) controls for the prefecture's AI stock one quarter prior to unrest events (and two quarters prior to the AI procurement outcome), and column (5) includes all controls. Panel A presents OLS regression estimates. Panel B presents a cross-fit partialing-out LASSO IV specification: we instrument for unrest events using weather variables interacted with themselves and an indicator for whether an unrest event occurred elsewhere in China on the day (variables are selected by LASSO), aggregated to the quarter level. All specifications include prefecture and quarter fixed effects. Standard errors are clustered by prefecture. * $p < .10$, ** $p < .05$, *** $p < .01$.

(column (3)), the past stock of public security AI procurement (column (4)), or all of these controls (column (5)).³³ One can see that across specifications, political unrest in a prefecture in one quarter is followed by a significantly greater amount of AI procurement in the following quarter.³⁴ The results remain qualitatively and quantitatively very similar throughout. The coefficients imply that a 1 standard deviation increase in local unrest is associated with around a 0.20 standard deviation increase in AI

33. As we show later, the lagged stock of AI procurement affects the occurrence of political unrest, mitigating the impact of weather variation on the occurrence of unrest; the lagged AI stock may also be correlated with subsequent AI procurement flows. It is thus important to control for the lagged AI stock as a potentially relevant omitted variable.

34. Our interpretation of AI procurement as a government response to political unrest suggests that firms receiving public security contracts issued following periods of political unrest should produce AI software for the government oriented toward surveillance. Indeed, we find a significant increase in the production of AI software intended for the government with surveillance functions (see [Online Appendix Figure A.XI](#) for details).

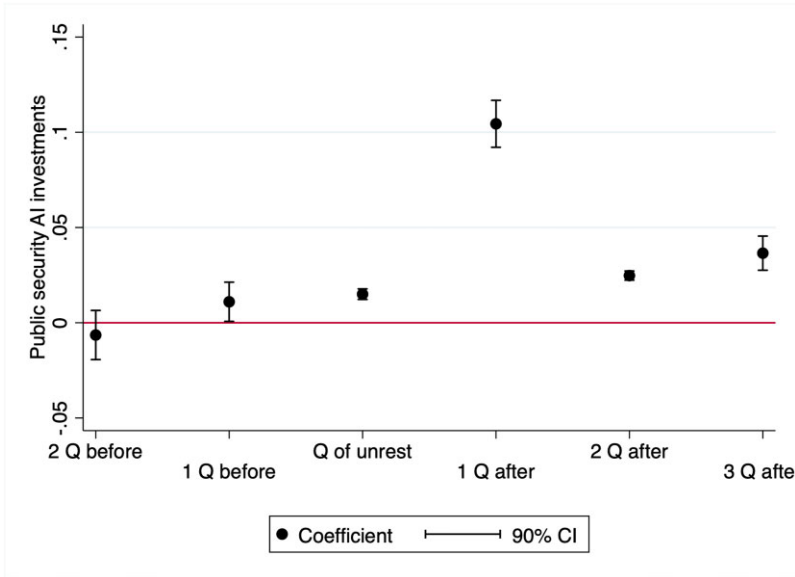


FIGURE II

Public Security AI Procurement Relative to the Quarter of Political Unrest

This figure plots the estimated effects of leads and lags of prefecture political unrest on prefecture public security AI procurement from a regression that also includes quarter and prefecture fixed effects. The outcome (AI procurement) and explanatory variables of interest (unrest events) are standardized to mean = 0, variance = 1.

procurement. [Online Appendix Table A.II](#) shows effects of political unrest by the separate subcategories of protests, public demands, and threats, with results remaining qualitatively the same. To the extent that reporting of these event types is subject to different degrees of censorship (e.g., due to differences in political sensitivity), these qualitatively similar patterns suggest that differential censorship of local unrest is unlikely to explain the baseline result.

We next examine whether AI procurement may already have been increasing in locations with political unrest prior to the unrest itself. We estimate a modified version of the baseline model, but also estimate the effects of unrest on AI procurement in periods from $t - 2$ to $t - 3$. [Figure II](#) plots the estimated coefficient on unrest for each lead and lagged period. As one can see, upcoming political unrest is associated with only slightly higher levels of AI procurement. The association between unrest and procurement

is substantially larger one quarter after the unrest occurrence, and such association fades in subsequent quarters. This pattern suggests that AI procurement primarily was in response to, and not in anticipation of, unrest.

As an alternative empirical strategy, we estimate the same panel model with locality and time fixed effects, now exploiting variation in unrest occurrence arising from daily weather conditions. Government officials may respond to occurrences of unrest even when they arise out of idiosyncratic weather shocks. This may be because officials are unable to distinguish between root causes of unrest, or because it is important to respond to any occurrence given the possible path dependence of unrest.

To implement an empirical strategy that instruments unrest occurrence using local weather conditions in our setting requires overcoming three challenges. The first challenge is high dimensionality: in a country as vast as China, one must consider a wide range of potentially relevant and interacting weather conditions. Moreover, in the Chinese context, weather conditions may affect unrest occurrence through multiple channels: most fundamentally, extreme weather conditions may stimulate political unrest due to socioeconomic hardship and potential unsatisfactory government responses. Weather conditions may affect the cost of political participation and the cost of police mobilization. This makes identifying a strong instrument more difficult and increases the researchers' degree of freedom and risk of finding false positives. The second challenge is the need to consider both the extensive and intensive margins of political unrest. Over a relatively long period of time, there are many days in which no unrest takes place (presumably because of the absence of mobilized political demands on those days), implying no elasticity between weather conditions and unrest occurrence. On certain days, unrest occurs across multiple prefectures, and local weather conditions plausibly would influence the likelihood of unrest occurrence in a specific location. A final challenge is the need to aggregate unrest occurrence to match the time frame over which AI procurement decisions are made (several months, which we operationalize as quarterly observations).

To address these challenges, we begin with the complete set of 18 weather variables consistently collected across weather stations in China. Reflecting the importance of weather interactions, we allow each variable to interact with the others. Reflecting the daily variation in the potential for (any) local unrest, we allow

the full set of weather variables to have heterogeneous effects on a prefecture's unrest occurrence depending on whether unrest occurs in at least one other prefecture on a given day. To identify a first stage while reducing the role of researcher discretion, we implement a LASSO regression to select predictors of unrest events among these weather variables, an indicator of unrest occurrence across China, and their interactions. Finally, we aggregate our first stage to the quarterly level and calculate the standard errors using the cross-fit partialing-out LASSO IV algorithm (Chernozhukov et al. 2018).³⁵

In Table II, Panel B, we present the estimated effect of unrest on AI procurement, now instrumenting for unrest using the LASSO IV.³⁶ We find that political unrest arising from local weather variation leads to significantly greater public security AI procurement in the subsequent quarter. This effect is robust to controlling for a variety of time-varying effects of local socioeconomic conditions. It is particularly noteworthy that controlling for the time-varying effects of local income and local government revenues—both of which might endogenously respond to variation in local weather (Dell, Jones, and Olken 2014)—does not affect our results. Nor does controlling for the lagged stock of public security AI that a locality has procured.³⁷ The IV analysis corroborates the OLS finding to provide further evidence that the relationship

35. Aggregating unrest events to the quarterly level matches the timing of procurement, and also addresses concerns about intertemporal substitution of unrest events within a narrow window of time.

36. In Online Appendix Table A.III, we present the weights assigned by LASSO to the selected weather predictors. The top three variables that LASSO selects are thunder \times unrest elsewhere, thunder \times visibility, and precipitation \times pressure \times unrest elsewhere. It is worth emphasizing that our IV strategy exploits the interaction of local weather and the potential for local unrest as measured by the occurrence of unrest elsewhere in China. We have examined whether weather conditions alone (without their interaction with unrest occurrence elsewhere in China) can predict public security AI procurement using the same LASSO specification, which can be thought of as a placebo reduced form. Importantly, we find that the LASSO specification identifies no weather variables to be predictive of procurement. Moreover, weather variables alone do not strongly predict local unrest either.

37. As we show later, a greater lagged stock of AI weakens the relationship between contemporaneous weather and unrest, thus our first stage is identified primarily from locations with lower AI stocks. We also find a weak effect of the lagged AI stock on subsequent unrest occurrence, independent of weather. We control for the lagged stock of AI in our IV specification as well (as shown in column (4)), and we find qualitatively and quantitatively similar effects.

between unrest and subsequent AI procurement is causal. We find that IV estimates are consistently larger than those of the OLS, potentially reflecting attenuation bias in the OLS estimates or the specific local average treatment effect estimated in the IV analysis.³⁸

Finally, as shown in [Figure III](#), we estimate very similar effects of unrest on AI procurement if we instead: (i) use a parsimonious set of weather conditions as first-stage predictors (rain, thunder, and wind speed); (ii) conduct our LASSO IV analysis as before, but measuring the potential for local unrest at a given time using the occurrence of political unrest in China within a week, rather than on the same day, to reduce measurement error in the first stage; and (iii) implement alternative estimation procedures using limited-information maximum likelihood (LIML) and jackknife IV estimators (JIVE). All estimates are positive and statistically significant at the 5% level, with magnitudes between the OLS estimate of around 0.20 and the baseline LASSO IV estimate just below 0.40.³⁹

1. *Upgraded Technology of Political Control.* Our evidence indicates a strong effect of political unrest on public security AI procurement. We provide further evidence that such procurement reflects an active decision by public security agencies to upgrade their technology of political control. Specifically, we examine whether these agencies make costly investments that could complement and enhance the efficacy of AI technology.

First, one would expect that the local government should invest in key hardware that complements facial recognition AI: high-resolution surveillance cameras, which provide the

38. The IV estimates may differ from the OLS estimates due to the particular characteristics of the compliers. We show below that prefectures with a low AI stock exhibit a weaker (subsequent) relationship between local weather and unrest. Such prefectures may be more likely to procure more AI in response to occurrence of local unrest, generating larger treatment effects. Of course one may have ex ante expected a relatively small estimate among the IV compliers, if one believed that weather-induced unrest reflects local political grievances that are marginal in nature, and hence would invite weaker responses by local governments in terms of subsequent investment in political control technology.

39. One may also be concerned about the robustness of the cross-fit partial-out LASSO algorithm, given the randomness in the process of drawing folds for cross-fitting. Across all specifications, we set the seed to one (only positive integers are available). Results using the first 100 seeds are shown under [Online Appendix Figure A.X](#). The 50th percentile coefficient estimate is 0.241.

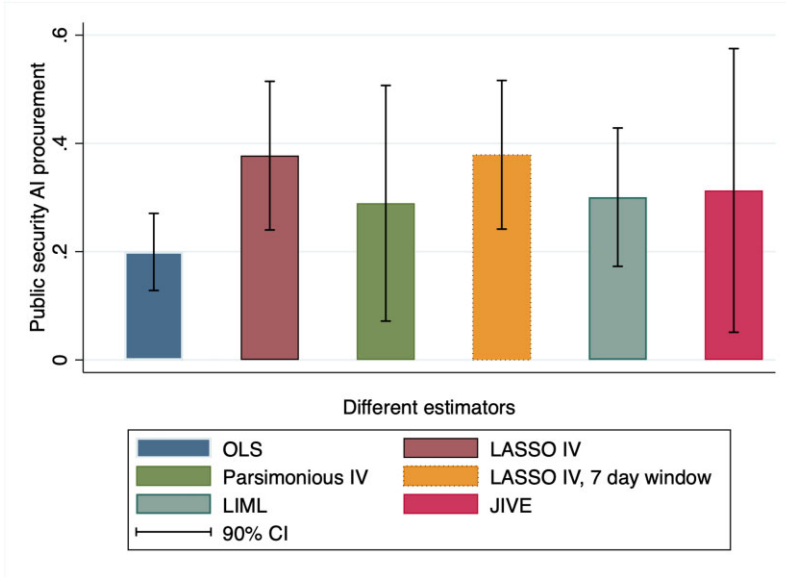


FIGURE III

Public Security AI Procurement with Different Estimators

Estimated effects of political unrest on public security AI procurement in the subsequent quarter (with 90% confidence intervals). The OLS bar replicates our baseline estimates (see Table II, Panel A, column (1)). The LASSO IV instruments for unrest using a cross-fit partialing-out LASSO IV algorithm on weather variables interacted with themselves and an indicator for whether unrest occurred elsewhere in China on the day (as in Table II, Panel B, column (1)). The parsimonious IV replicates this specification using a more parsimonious set of weather variables (interacting rain, thunder, and wind with unrest elsewhere in China on the day). The LASSO IV, seven-day window, expands the first-stage window for unrest elsewhere in China to one week instead of limiting it to the same day. LIML and JIVE replicate the same specification using these alternate estimators, including the same instruments used by LASSO. All specifications include prefecture and quarter fixed effects. The outcome and independent variables are standardized to mean = 0, variance = 1.

fundamental video data processed by the AI algorithm. In Table III, Panels A and B, we replicate the exercise in Table II but examine the local public security procurement of surveillance cameras. We find that following the occurrence of political unrest, the local public security units also increase their procurement of high-resolution surveillance cameras. The timing of surveillance cameras' procurement matches that of the AI procurement, with a substantial increase in the quarter after the unrest occurrence

TABLE III
THE EFFECT OF UNREST EVENTS ON SURVEILLANCE CAMERA AND
FACIAL-RECOGNITION AI PROCUREMENT

	Public security AI/camera procurement				
	(1)	(2)	(3)	(4)	(5)
Panel A: OLS, cameras					
Unrest events _{<i>t</i>-1}	0.436*** (0.084)	0.420*** (0.083)	0.436*** (0.085)	0.425*** (0.080)	0.425*** (0.080)
Panel B: LASSO IV, cameras					
Unrest events _{<i>t</i>-1}	0.593*** (0.175)	0.599*** (0.174)	0.593*** (0.175)	0.559*** (0.167)	0.559*** (0.167)
Panel C: OLS, AI × surveillance cameras					
Unrest events _{<i>t</i>-1}	0.681*** (0.154)	0.669*** (0.157)	0.680*** (0.155)	0.671*** (0.148)	0.671*** (0.148)
Panel D: Lasso IV, AI × surveillance cameras					
Unrest events _{<i>t</i>-1}	1.054*** (0.374)	1.070*** (0.376)	1.054*** (0.374)	0.967*** (0.334)	0.966*** (0.334)
GDP × quarter	Yes	No	No	No	Yes
Log population × quarter	No	Yes	No	No	Yes
Gov. revenue × quarter	No	No	Yes	No	Yes
AI stock _{<i>t</i>-2}	No	No	No	Yes	Yes

Notes. This table presents regressions at the prefecture-quarter level. The outcome in Panels A and B is the number of surveillance cameras procured by the local government. The outcome in Panels C and D is the product of the number of public security facial-recognition AI contracts procured by the local government and the number of surveillance cameras procured by the corresponding government. Outcome variables in all panels are standardized to mean = 0 and variance = 1. The explanatory variable of interest is the occurrence of unrest events in the corresponding prefecture during the preceding quarter, also standardized. Column (1) controls for prefecture GDP × quarter effects, column (2) controls for log prefecture population × quarter effects, column (3) controls for prefectural government tax revenue × quarter effects, column (4) controls for the prefecture's AI stock one quarter prior to unrest events (and two quarters prior to the AI procurement outcome), and column (5) includes all controls. Panels A and C present OLS regression estimates. Panels B and D present a cross-fit partialing-out LASSO IV specification: we instrument for unrest events using weather variables interacted with themselves and an indicator for whether an unrest event occurred elsewhere in China on the day (variables are selected by LASSO), aggregated to the quarter level. All specifications include prefecture and quarter fixed effects. Standard errors are clustered by prefecture. * $p < .10$, ** $p < .05$, *** $p < .01$.

(see [Online Appendix Figure A.XII](#)). Examining the locality's decision to jointly procure AI technologies and surveillance cameras, measured as the product of the two, we find a similar (but larger in magnitude) effect, reflecting public security agencies' decision to invest in both following unrest occurrence (see [Table III](#), Panels C and D, and [Online Appendix Figure A.XIII](#)).

Second, one may expect changes in the local public security agencies' personnel arrangements as they increasingly deploy AI technologies ([Acemoglu et al. 2022](#)). In particular, it has been argued by [Acemoglu and Restrepo \(2020\)](#) and

Agrawal, Gans, and Goldfarb (2019) that AI technology is often labor-saving and likely to be skill-biased. In the context of public security agencies, AI technology may substitute for patrol officers while still necessitating desk officers to analyze the AI output. Consistent with these predictions, we find that local police hiring is significantly lower one year after the corresponding police department procures AI technology, and the share of desk (as opposed to patrol) police significantly increases among the new hires (see [Online Appendix Table A.IV](#) for details). This suggests that the local public security agencies adjust their personnel composition alongside the deployment of facial-recognition AI.

Taken together, these results suggest that the autocrat and their public security arms view AI technology as potentially useful and actively procure AI as an advanced method of political control. Moreover, the increased procurement of AI represents a component of a coherent technological bundle along with high-resolution surveillance cameras and relatively skilled labor in the police force, which could complement AI and help the autocrat maintain political control.

III.B. The Effect of AI Procurement on the Occurrence of Unrest

We examine whether greater AI procurement by the local governments' public security agencies effectively suppresses political unrest. Anecdotally, local governments appear to deploy facial-recognition AI to reduce unrest through means such as identifying new faces in a protest, tracking suspicious persons in their daily lives, or deterring potential participants in unrest.⁴⁰

Importantly, having just demonstrated that AI procurement is endogenous to political unrest, we cannot directly estimate the effect of it on subsequent political unrest. Estimating such a relationship is further challenged by the potential for strong autocorrelation over time in local political unrest.⁴¹

To evaluate the effect of public security AI procurement on autocrats' political control, we thus examine how past public security AI procurement shapes the effects of external shocks on local political unrest. Consider a context in which multiple locations share a common elevated potential for political unrest but

40. For example, see [Andersen \(2020\)](#).

41. Unrest occurrence in a given quarter is strongly, positively associated with unrest in the subsequent quarter. The estimated autocorrelation coefficient is 0.30 (p -value < .001).

experience different idiosyncratic weather conditions that shape the occurrence of unrest (as we have demonstrated in the LASSO IV first stage in the analysis above). In such a context, one may wonder whether the preexisting stock of AI technology procured by the public security agencies may mitigate the effect of weather conditions that are generally conducive to the occurrence of political unrest.

To test this hypothesis, we estimate the effects of contemporaneous weather shocks in prefecture i at time t on local political unrest, allowing this effect to vary depending on the lagged stock of local public security procurement of AI up to period $t - 1$, controlling for prefecture and time period (quarter) fixed effects. Specifically, we estimate the following model:

$$\begin{aligned}
 \text{Unrest}_{it} = & \beta_1 \text{AI_stock}_{i,t-1} + \beta_2 \text{ConduciveWeather}_{it} \\
 & + \beta_3 \text{ConduciveWeather}_{it} \times \text{AI_stock}_{i,t-1} + \alpha_t + \gamma_i \\
 (2) \quad & \delta_t X_i + \epsilon_{it}
 \end{aligned}$$

Table IV, Panel A, column (1), presents the baseline result. As we saw already, weather conducive to political unrest is positively and significantly associated with the occurrence of unrest (as expected from our LASSO regression results). However, the estimated effect of conducive weather interacted with the stock of public security AI procurement is negative: past accumulation of AI capacity significantly weakens the positive relationship between conducive weather and unrest occurrence, suggesting a role of AI in maintaining political control. In other words, when weather conditions are favorable and localities exhibit some potential for unrest, such unrest is more likely to occur. However, this unrest is significantly less likely to take place in localities with greater stocks of public security AI procurement. Figure IV presents the results visually, where we plot the relationship between local weather conditions on contemporaneous unrest occurrence, across localities with a high level of lagged stock of public security AI procurement (two standard deviations above the mean) and a low level (two standard deviations below the mean). One observes that although the past stock of AI does not substantially change unrest occurrence in below-average weather conditions, the responsiveness of unrest occurrence to above-average

TABLE IV
THE EFFECT OF AI PROCUREMENT ON THE OCCURRENCE OF UNREST

	Standardized number of unrest events			
	(1)	(2)	(3)	(4)
Panel A: Effect of stock of public security AI				
Conducive weather	0.9176*** (0.1609)	0.9523*** (0.1597)	0.9183*** (0.1613)	0.9511*** (0.1543)
Public security procurement stock AI_{t-1}	-0.0080** (0.0039)	-0.0032 (0.0050)	-0.0079** (0.0038)	-0.0020 (0.0050)
Conducive weather \times public security AI_{t-1}	-0.2265* (0.1153)	-0.2729** (0.1306)	-0.2260* (0.1156)	-0.2662** (0.1250)
Panel B: Effect of stock of cameras \times public security AI				
Conducive weather	0.9113*** (0.1585)	0.9446*** (0.1560)	0.9118*** (0.1587)	0.9449*** (0.1517)
Public security procurement stock cam. and AI_{t-1}	0.2462** (0.1074)	0.2734*** (0.0997)	0.2455** (0.1073)	0.2638*** (0.0945)
Conducive weather \times public security cam. and AI_{t-1}	-0.5688** (0.2281)	-0.6598*** (0.2401)	-0.5735** (0.2304)	-0.6403*** (0.2229)
GDP \times quarter	Yes	No	No	Yes
Log population \times quarter	No	Yes	No	Yes
Gov. revenue \times quarter	No	No	Yes	Yes

Notes. This table presents regressions at the prefecture-quarter level. The outcome of interest is the number of political unrest events in the prefecture in a given quarter, standardized to mean = 0 and variance = 1. Conducive weather is the standardized number of predicted unrest events (aggregated to the quarter level) from the LASSO specification discussed in the text. The stock of public security AI and surveillance camera procurement are also standardized to mean = 0 and variance = 1. Column (1) controls for prefecture GDP \times quarter fixed effects, column (2) controls for log prefecture population \times quarter fixed effects, column (3) controls for prefectural government tax revenue \times quarter fixed effects, and column (4) includes all prior controls. All specifications include prefecture and quarter fixed effects. Standard errors are clustered by prefecture. * $p < .10$, ** $p < .05$, *** $p < .01$.

conditions diverges.⁴² We continue to find qualitatively and quantitatively similar results as we gradually add time-varying controls to account for changes in local socioeconomic conditions (shown in Table IV, Panel A, columns (2)–(4)). A one standard deviation increase in the stock of past public security AI procurement

42. It is important to note that in our model with linear interactions, the lagged AI stock is predicted to increase unrest when the weather is not conducive to unrest. But as can be seen in Figure IV, this estimated differential positive effect of AI stocks is not statistically significant.

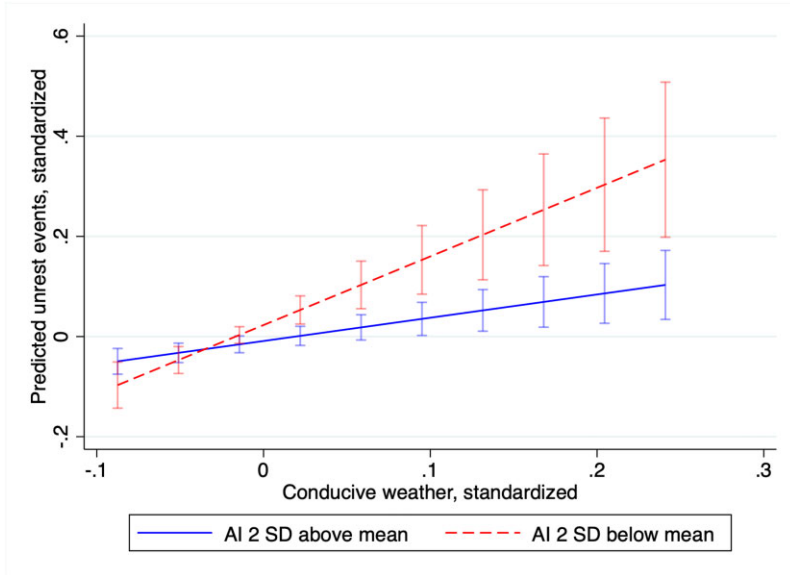


FIGURE IV

The Effects of Weather Conducive to Political Unrest

The effect of weather conducive to political unrest varying according to the lagged stock of prefecture public security AI. This figure displays the predicted effect of conducive weather (from the LASSO specification) on the number of political unrest events in the prefecture, at two levels of the stock of AI in the quarter before unrest: at two standard deviations above the mean (solid blue line) and at two standard deviations below the mean (dashed red line) (color version available online). All specifications include prefecture and quarter fixed effects. The outcome and independent variables are standardized to mean = 0, variance = 1. Along the x -axis, conducive weather ranges from the 5th percentile to the 95th percentile of the data.

cuts by one-quarter the effect of weather conditions conducive to political unrest.⁴³

Next we examine whether complementary technological investment increases the effectiveness of facial-recognition AI in suppressing unrest. We estimate the baseline specification,

43. We again find qualitatively similar results for each subcategory of the unrest events (protests, public demands, and threats); see [Online Appendix Table A.V](#). To the extent that these distinct event types are subject to different degrees of censorship in reporting of local unrest, this suggests that the results we find are unlikely to be explained by confounding factors that are correlated with both local governments' procurement of facial recognition AI technology and its use of censorship.

replacing $AI_stock_{i,t-1}$ with the interaction between the lagged stock of AI and the stock of cameras $AI_stock_{i,t-1} \times cameras_stock_{i,t-1}$ (shown in Table IV, Panel B). We observe that when cameras are procured alongside facial-recognition AI, the effectiveness of AI becomes amplified. A one standard deviation increase in the joint procurement stock of AI and surveillance cameras cuts the effect of conducive weather in half, an even larger effect than that of AI alone.⁴⁴

This empirical strategy relies on plausibly exogenous variation in weather conditions but endogenous variation in the lagged stock of AI procurement. We examine two sets of potential confounding variables related to the stock of AI—one regarding local governance and capacity more broadly and the other concerning past unrest and changing local political dynamics. First, one might worry that past AI procurement for any purpose reflects local governments embracing new technology and more broadly the quality of local governance, which may in turn dampen political unrest. To address this concern, we conduct a placebo test: does local government's AI procurement outside the public security agencies in the past shape the relationship between local weather conditions and unrest occurrence? Crucially, the effect of past AI procurement only appears for the contracts issued by public security agencies. Local AI procurement by non-public security agencies does not mitigate the effects of conducive weather on political unrest, as shown in Table V, Panel A.⁴⁵ Second, because the cross-prefecture variation in previous AI procurement is partially shaped by past political unrest (as shown above), if the past unrest is associated

44. Interestingly, when we consider only the effect of the lagged stock of surveillance cameras $camera_stock_{i,t-1}$, we find that the stock of high-resolution surveillance cameras alone does not effectively suppress subsequent unrest occurrence (see Online Appendix Table A.VI).

45. Relatedly, one may also be concerned that deployment of facial-recognition AI in response to unrest captures local politicians' strong career incentives, which could be associated with a range of other policies also aimed at suppressing subsequent unrest. To assess this possibility, we examine whether exogenous weather shocks have heterogeneous effects on contemporaneous unrest occurrence depending on local politicians' career incentives. We follow Wang, Zhang, and Zhou (2020) and estimate an index capturing each prefectural city leader's ex ante likelihood of promotion in each year, as a flexible function of their age (relative to retirement), tenure, and official rank in the bureaucratic system (capturing the potential for upward mobility). As shown in Online Appendix Table A.VII, we do not find a noticeable pattern of heterogeneous effects of conducive weather depending on local politicians' career incentives.

TABLE V
 FALSIFICATION EXERCISES: THE EFFECTS OF NON-PUBLIC SECURITY AI PROCUREMENT
 AND PAST UNREST

	Standardized number of unrest events			
	(1)	(2)	(3)	(4)
Panel A: Effect of stock of non-public security AI				
Conductive weather	0.9378*** (0.1678)	0.9768*** (0.1666)	0.9385*** (0.1682)	0.9747*** (0.1608)
Non-public security procurement stock AI _{t-1}	-0.0021* (0.0012)	-0.0022 (0.0014)	-0.0021* (0.0012)	-0.0019 (0.0012)
Conductive weather × non-public security AI _{t-1}	-0.0441 (0.0299)	-0.0513 (0.0338)	-0.0444 (0.0305)	-0.0473 (0.0301)
Panel B: Effect of past unrest				
Conductive weather	0.9770*** (0.1660)	0.9813*** (0.1711)	0.9777*** (0.1664)	0.9798*** (0.1656)
Unrest _{t-1}	0.0002 (0.0672)	-0.0006 (0.0706)	0.0038 (0.0684)	0.0013 (0.0664)
Conductive weather × unrest _{t-1}	-0.0110 (0.0187)	-0.0114 (0.0195)	-0.0120 (0.0189)	-0.0114 (0.0186)
GDP × quarter	Yes	No	No	Yes
Log population × quarter	No	Yes	No	Yes
Gov. revenue × quarter	No	No	Yes	Yes

Notes. This table follows Table IV and presents regressions at the prefecture-quarter level. The outcome of interest is the number of political unrest events in the prefecture in a given quarter, standardized to mean = 0 and variance = 1. Conductive weather is the standardized number of predicted unrest events (aggregated to the quarter level) from the LASSO specification discussed in the text. The stock of non-public security AI and local unrest prior to the unrest event outcome are also standardized to mean = 0 and variance = 1. Column (1) controls for prefecture GDP × quarter fixed effects, column (2) controls for log prefecture population × quarter fixed effects, column (3) controls for prefectural government tax revenue × quarter fixed effects, and column (4) includes all prior controls. All specifications include prefecture and quarter fixed effects. Standard errors are clustered by prefecture. * $p < .10$, ** $p < .05$, *** $p < .01$.

with heterogeneity in the responsiveness of unrest to subsequent weather shocks, this could confound the interpretation that results capture the effects of public security AI procurement.⁴⁶ To assess this possibility, we examine whether exogenous weather

46. For example, past unrest may make subsequent unrest more likely (e.g., due to path dependence), which could potentially increase the elasticity of unrest occurrence with respect to contemporaneous weather conditions. Alternatively, past unrest may reduce the likelihood of subsequent unrest (e.g., due to increased overall government repression independent of AI), thus reducing the elasticity of unrest occurrence with respect to weather conditions.

shocks have heterogeneous effects on contemporaneous unrest occurrence depending on past unrest in the locality. Specifically, we estimate specifications analogous to those described above, replacing $AI_stock_{i,t-1}$ with $unrest_{i,t-1}$. As shown in Table V, Panel B, we do not find a noticeable pattern of heterogeneous effects of conducive weather depending on past unrest in the locality, suggesting that the pattern of heterogeneity we observe is likely due to public security AI procurement, rather than other mechanisms arising from past unrest per se.⁴⁷

Taken together, these results suggest that the politically motivated procurement of AI technology is indeed useful in enhancing the state's political control capacity.

IV. THE ROLE OF AUTOCRATIC POLITICAL CONTROL IN AI INNOVATION

We turn to the question of whether politically motivated procurement of AI stimulates AI innovation. Specifically, we focus on AI procurement contracts issued by public security agencies in prefectures that experienced above-median levels of political unrest in the quarter prior to the contracts' issuance. As shown in the previous section, these contracts are plausibly issued for purposes of political control.

We use a staggered event study design to identify the overall effects of procurement contracts issued for political control on the subsequent product development and innovation among the facial-recognition AI firms that are awarded the contracts. The empirical strategy exploits variation across time and across firms in the receipt of a government contract.

As in an event study design, we compare firms' outcomes—their software releases—before and after they receive their first politically motivated public security contracts, controlling for firm and time period fixed effects.⁴⁸ Specifically, we estimate the effect of firms receiving their first government contracts when these are

47. We are also able to directly examine one margin along which local governments may respond to past unrest: police hiring. We estimate the baseline specification (as in Table IV, Panel A, column (4)), also controlling for past police hiring and its interaction with contemporaneous weather conditions. We find that the mitigating effect of the lagged public security AI stock is nearly unchanged: the coefficient on $ConduciveWeather_{it} \times AI_stock_{i,t-1}$ is -0.2718 (p -value = $.0338$).

48. We only examine firms' first contracts because subsequent contracts could be endogenous to firms' performance in the initial contracts.

public security contracts issued in a prefecture that recently experienced political unrest. We estimate the following specification:

$$(3) \quad y_{it} = \beta_{1T} T_{it} + \alpha_t + \gamma_i + \delta_t X_i + \epsilon_{it},$$

where T_{it} equals 1 if, at time t , T quarters have passed before/since firm i received its first politically motivated public security contract; α_t are a full set of quarter fixed effects; and γ_i are a full set of firm fixed effects. The coefficients β_{1T} describe software production of a firm around the time when it receives its first politically motivated public security procurement contract.

In [Figure V](#), we plot the series of β_{1T} coefficients, considering the cumulative, total software output as well as output for government and commercial applications, respectively. In Panel A, one can see that firms receiving a public security contract issued following episodes of political unrest develop approximately 10 more software products over the next two years, representing an increase in software of around 30% of one standard deviation. One naturally wonders whether firms receiving these contracts were already following a different trend of software production before the receipt of the contracts. However, conditional on firm fixed effects, we do not observe differential precontract software production levels or trends among firms that would go on to receive a public security procurement contract. We present event study coefficients from cumulative software production eight quarters after contract in [Table VI](#), column (1); we present the coefficients from a specification where we control for time-varying effects of an index that captures firms' underlying potential to benefit from a contract, composed of firms' precontract size and contract value, in column (2); and we present coefficients from a weighted event study specification, following [Borusyak, Jaravel, and Spiess \(2017\)](#), in column (3).⁴⁹ The full set of event study coefficients are presented in [Online Appendix Table A.VIII](#).

In [Figure V](#), Panels B and C, we separately present results for software products intended for government and for commercial purposes, respectively. One observes that firms receiving

49. The inclusion of the financial value of the contract as a control variable shuts down one mechanism through which politically motivated contracts shape innovation and is thus arguably overcontrolling for unobserved drivers of firms' selection into contracts.

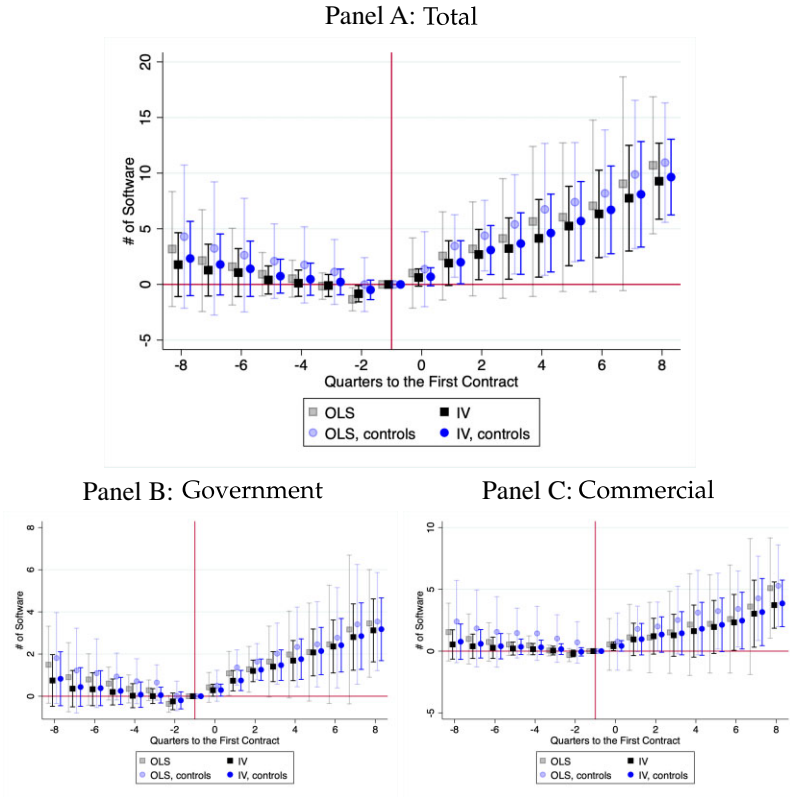


FIGURE V

The Effects of Politically Motivated Contracts

The figure shows the effects of first contracts for facial-recognition AI firms that earn contracts from public security arms of local governments when there is an above-median amount of unrest in the quarter prior to the contract. The figure shows the total software production (Panel A), software developed for government applications (Panel B), and software developed for commercial applications (Panel C), relative to the time of receiving the initial contract. All estimates control for firm and time period fixed effects. Gray lines/square markers show the baseline estimated effect over time for firms. Blue lines/circle markers add controls for the time-varying effects of an index that captures firms' underlying potential to benefit from a contract, composed of firms' precontract size and contract value. Dark lines/markers use LASSO-selected weather variables to instrument for unrest.

public security procurement contracts following episodes of political unrest not only differentially produce more software for the government, but also increase their commercial software development. The differential increase in commercial software

TABLE VI
THE TOTAL EFFECT OF POLITICALLY MOTIVATED PUBLIC SECURITY CONTRACTS ON SOFTWARE PRODUCTION

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Total software						
8 quarters after contract	10.671*** (3.664)	11.153*** (3.103)	13.111*** (2.957)	9.287*** (2.127)	9.711*** (2.119)	11.124*** (1.948)
Panel B: Government software						
8 quarters after contract	3.465** (1.543)	3.566** (1.415)	4.198*** (1.375)	3.099*** (0.920)	3.170*** (0.916)	3.745*** (0.952)
Panel C: Commercial software						
8 quarters after contract	5.098** (2.409)	5.310** (1.956)	6.009*** (1.856)	3.718*** (1.146)	3.898*** (1.157)	4.373*** (0.945)
Unrest events	All	All	All	Instrumented	Instrumented	Instrumented
Firm characteristics	No	Yes	No	No	Yes	No
Event study weighting	No	No	Yes	No	No	Yes

Notes. The table presents regression coefficients for facial-recognition AI firms that earn contracts from local governments when there is an above-median amount of unrest in the quarter before the contract. The table shows the difference in total software production between firms that earn politically motivated (public security) contracts versus firms that do not earn a contract. Panel A uses total software production as the outcome, Panel B uses government software, and Panel C uses commercial software. Columns (1)–(3) measure unrest using all observed events. Columns (4)–(6) measure unrest using the predicted unrest based on the LASSO IV specification discussed in the text. All columns control for time period fixed effects and firm fixed effects. Columns (2) and (5) include controls for the time-varying effects of the contract and company size (an inverse covariance-weighted z-score for contract size and company size interacted with year indicators, following [Andersson 2008](#)). Columns (3) and (6) weight the control group 1,000 times more than the treatment, following [Boruyak, Jaravel, and Spiess 2017](#). The full set of coefficients can be found in [Online Appendix Tables A.VIII–A.X](#). Standard errors are clustered at the contract location (prefecture) level. * $p < .10$, ** $p < .05$, *** $p < .01$.

development totals around five more software products over two years after the contract receipt, representing an increase of around 30% of a standard deviation.⁵⁰ Again we observe no differential software production level or trend for either government or commercial categories prior to the receipt of the public security contracts, suggesting a causal interpretation. Our findings indicate a role for politically motivated government procurement in stimulating frontier innovation for both government and commercial applications.⁵¹

One concern with this analysis is that our definition of politically motivated contracts relies on the endogenous occurrence of political unrest. Factors that shape political unrest may be associated with production of AI software specifically among firms that select into public security contracts (though recall that they are not generally local firms, so time- and location-varying shocks do not directly drive these results). To address this concern, we alternatively define a politically motivated contract as a public security contract issued just after a quarter with above-median predicted political unrest, using our weather-based LASSO instruments to predict unrest, as described in [Section III](#). The estimated coefficients from this alternative definition of politically motivated contracts are plotted in darker-shaded dots in [Figure V](#), and presented in [Table VI](#), columns (4)–(6) (see [Online Appendix Table A.VIII](#), columns (4)–(6), for the full set of event study coefficients). One can see that the effects of public security contracts on software innovation are very similar following episodes of plausibly exogenous political unrest.

Another concern is that the baseline specification may be capturing the effects of mechanisms other than the politically

50. We present the full set of event study coefficients for commercial and government software in [Online Appendix Tables A.IX and A.X](#).

51. As an auxiliary test of the role of access to large quantities of government data collected for political motivations, we examine whether firms receiving public security contracts in a politically sensitive environment develop data-complementary tools (e.g., software supporting data storage) to manage the large quantities of data that they receive access to. Importantly, these data-complementary software products are distinct from the AI software studied here. In [Online Appendix Figure A.XIV](#), we present estimates from the same specification as in [Figure V](#), but now considering the outcome of data-complementary software products. One can see that data-complementary software production differentially increases after the receipt of a public security contract in a politically sensitive environment, relative to the receipt of a non-public security contract.

motivated public security contract per se. For example, firms receiving contracts in politically sensitive contexts (i.e., just following episodes of local unrest) may be specially selected in a way that may be related to subsequent performance. These firms may also develop political connections (needed to receive a contract at a politically sensitive moment), which might affect later performance. To address this concern, we compare the effects of public security contracts issued in a politically sensitive environment (defined as municipalities with above-median political unrest in the previous quarter) with those of non-public security contracts issued in the same environment.⁵² Specifically, among firms receiving their first government contracts in a prefecture that recently experienced political unrest, we estimate the following specification:

$$(4) \quad y_{it} = \beta_{1T} T_{it} + \beta_{2T} T_{it} \times PublicSecurity_i + \alpha_t + \gamma_i + \delta_t X_i + \epsilon_{it},$$

where $PublicSecurity_i$ is an indicator that the firm's first government contract is issued by a public security agency. The coefficients β_{1T} describe software production of a firm around the time when it receives its first government contract when this contract is issued by a non-public security agency (in a politically sensitive context); the sums of coefficients $\beta_{1T} + \beta_{2T}$ describe software production around the time when a firm receives its first government contract when this contract is issued by a public security agency; and the sequence of coefficients β_{2T} thus captures the differential software production before and after a firm receives the contract in a politically sensitive environment.

In [Figure VI](#), we plot the series of β_{2T} coefficients; in [Table VII](#), we present the regression estimates.⁵³ Compared with firms receiving a non-public security contract issued in the same local political environment, we continue to observe a positive and significant effect of politically motivated public security contracts on firms' total software production over the next two years, as

52. The classification of a public security contract is only dependent on the agency issuing the contract; thus, public security contracts can be categorized as non-politically motivated if they were issued in a prefecture that had not experienced above-median unrest in the previous quarter.

53. In [Online Appendix Tables A.XI–A.XIII](#), we present the full set of event study coefficients.

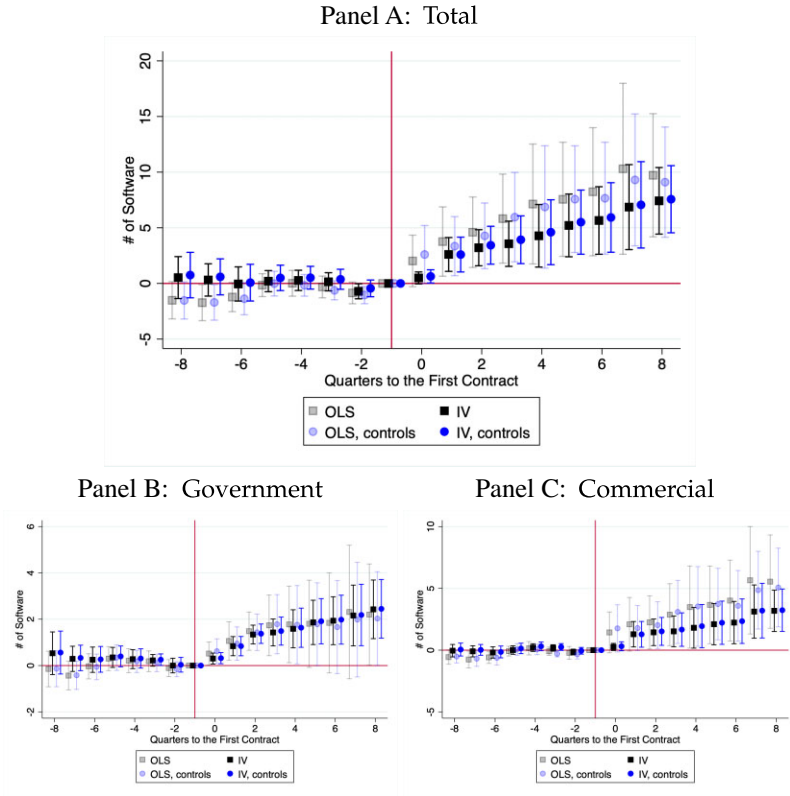


FIGURE VI

Differential Effects of Politically Motivated Contracts

Effects of politically motivated contracts on total software development (Panel A), software developed for government applications (Panel B), and software developed for commercial applications (Panel C). All panels restrict firms to those that receive their first contracts in prefectures experiencing above-median political unrest (or predicted unrest) in the previous quarter and control for firm and time period fixed effects. Panels compare firms receiving public security contracts to those that receive contracts from other agencies. Gray lines/square markers show the baseline estimated effect over time for firms. Blue lines/circle markers add controls for the time-varying effects of an index that captures firms' underlying potential to benefit from a contract, composed of firms' precontract size and contract value. Dark lines/markers use LASSO-selected weather variables to instrument for unrest.

TABLE VII
THE DIFFERENTIAL EFFECT OF POLITICALLY MOTIVATED PUBLIC SECURITY CONTRACTS ON TOTAL SOFTWARE PRODUCTION

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Total software						
8 quarters before contract	5.239 (3.427)	4.170 (2.805)	0.646 (1.318)	1.509 (1.539)	0.564 (1.519)	-1.043 (0.926)
8 quarters after contract	0.823 (1.644)	4.625** (1.946)	6.330*** (1.696)	1.893* (1.029)	4.041*** (1.322)	5.452*** (1.285)
8 quarters before contract \times public security	-1.897 (1.239)	-2.221 (1.665)	-1.941 (1.221)	0.246 (1.250)	0.464 (1.453)	-0.055 (1.286)
8 quarters after contract \times public security	9.825*** (3.368)	4.024 (3.829)	7.896*** (2.439)	7.460*** (1.863)	4.071** (1.670)	6.385*** (1.560)
Panel B: Government software						
8 quarters before contract	1.697 (1.067)	0.951 (0.706)	0.065 (0.514)	0.230 (0.491)	-0.259 (0.631)	-0.712 (0.508)
8 quarters after contract	1.278 (0.778)	3.196*** (1.054)	3.205*** (0.885)	0.700 (0.490)	1.465** (0.702)	2.048*** (0.639)
8 quarters before contract \times public security	-0.158 (0.463)	0.154 (0.602)	-0.222 (0.457)	0.429 (0.498)	0.883 (0.643)	0.250 (0.477)
8 quarters after contract \times public security	2.180* (1.284)	-0.128 (1.455)	1.360 (1.079)	2.408*** (0.771)	1.453 (0.906)	1.963*** (0.756)

TABLE VII
CONTINUED

	(1)	(2)	(3)	(4)	(5)	(6)
Panel C: Commercial software						
8 quarters before contract	2.267 (1.510)	1.694 (1.015)	0.304 (0.490)	0.689 (0.742)	0.285 (0.555)	-0.263 (0.276)
8 quarters after contract	-0.420 (0.821)	0.759 (0.637)	1.883** (0.703)	0.476 (0.551)	1.174*** (0.410)	1.820*** (0.456)
8 quarters before contract × public security	-0.754 (0.511)	-0.937* (0.550)	-0.823 (0.563)	-0.171 (0.360)	-0.169 (0.393)	-0.274 (0.415)
8 quarters after contract × public security	5.482** (2.168)	1.979 (1.616)	4.521*** (1.522)	3.225*** (1.037)	1.552*** (0.636)	2.805*** (0.803)
Unrest events	All	All	All	Instrumented	Instrumented	Instrumented
Firm characteristics	No	Yes	No	No	Yes	No
Event study weighting	No	No	Yes	No	No	Yes

Notes. The table presents regression coefficients for facial-recognition AI firms that earn contracts from local governments when there is an above-median amount of unrest in the quarter prior to the contract. The table shows the difference in software production between firms that earn politically motivated (public security) contracts versus non-politically motivated (non-public security) contracts. Panel A uses total software production as the outcome, Panel B uses government software, and Panel C uses commercial software. Columns (1)–(3) measure unrest using all observed events. Columns (4)–(6) measure unrest using the predicted unrest events based on the LASSO IV specification discussed in the text. All columns control for time period fixed effects and firm fixed effects. Columns (2) and (5) include controls for the time-varying effects of the contract and company size (an inverse covariance-weighted z-score for contract size and company size interacted with year indicators, following Anderson 2008), Columns (3) and (6) weight the control group 1,000 times more than the treatment, following Borusyak, Jaravel, and Spiess (2017). The full set of coefficients can be found in Online Appendix Tables A.XI–A.XIII. * $p < .10$, ** $p < .05$, *** $p < .01$.

well as software intended for government and commercial purposes. We do not observe differential precontract software production levels or trends among firms that went on to receive a public security procurement contract in a politically sensitive environment. Defining politically sensitive environments using the predicted level of political unrest based on our weather-based LASSO instruments (rather than observed unrest) yields very similar results.

IV.A. Robustness and Ruling Out Alternative Hypotheses

The results presented thus far do not appear to be the result of differential selection by firms into politically motivated public security procurement contracts. Although we cannot absolutely rule out some role for differential selection of firms into contracts, we find no evidence of precontract differences in software production levels or trends, which one would expect if firms selected into these contracts as a function of their underlying productivity. As an additional check, we flexibly control for the time-varying effects of firms' age and precontract software production to address concerns about firms selecting into contracts as a function of their potential production growth (see [Online Appendix](#) Tables A.XIV–A.XV, Panels A.2 and A.3). The results are qualitatively and quantitatively similar across these alternative specifications.

We assess the robustness of our results to variation in specifying our outcome of interest—measures of software innovation. We restrict attention only to firms' new software releases (i.e., version 1.0) and major upgrades with a change in the first digit of the release number (i.e., versions 2.0, 3.0). Our baseline estimates remain largely unchanged, indicating that our results are not driven by minor software updates (see [Online Appendix](#) Tables XIV–XV Panel B). Moreover, we consider software in a field of AI that is considered most difficult and frontier in its application: video-based facial recognition, which (as opposed to static images) requires N -to-1 or even N -to- N matching algorithms, and we find qualitatively similar results (see [Online Appendix](#) Tables XIV–XV Panel C).

Given the complex process of constructing our data set, it is important to note that our findings are robust to varying several salient dimensions of our analysis (see [Online Appendix](#) Tables A.XIV–A.XV, Panels D–F). First, our results are robust to adjusting our classification of public security contracts to exclude any government contracts ambiguously related to public

security (e.g., contracts with the government headquarters and smart city management and administrative bureaus could be meant to provide security services just for the government office building; see [Online Appendix](#) Tables XIV–XV Panel D). Second, the results are robust to adjustments of the parameters of the machine-learning algorithm used to classify software—timestep, embedding, and nodes of the RNN LSTM model (see [Online Appendix](#) Tables XIV–XV Panel E). Third, our results are robust to considering a balanced panel of firms within a narrow window, and to expanding the window of time around the receipt of the first contract that we study (see [Online Appendix](#) Tables XIV–XV Panel F).

Our results are also maintained under specifications that help us address a range of alternative hypotheses. One concern is that contracts with the public security agencies in the powerful, high-surveillance local governments of Beijing or Shanghai may not generalize to the broader range of politically motivated contracts. To rule out the possibility that our findings are distorted by contracts with these two local governments, we estimate our baseline specification and add fixed effects for contracts from Beijing and Shanghai governments interacted with a full set of quarter to/from contract fixed effects (see [Online Appendix](#) Tables XIV–XV Panel G.1). Results are also robust to dropping contracts from the surveillance-intensive province of Xinjiang (see [Online Appendix](#) Tables XIV–XV Panel G.2). We also account for a firm's home prefecture/province government potentially giving the firm a commercial advantage beyond the procurement contracts themselves by estimating the baseline model excluding contracts signed between firms and any government in their home prefecture/province (see [Online Appendix](#) Tables XIV–XV Panels G.3 and G.4). Finally, to address a broader set of concerns about time- and space-varying shocks that may drive firms' commercial activities, we control for province by quarter fixed effects and show that results are qualitatively similar (see [Online Appendix](#) Tables XIV–XV Panel H).

IV.B. Software Export Activities

Firms' export activities are often considered a signal of production at the technological frontier ([Vernon 1966](#); [Melitz 2003](#); [Filatotchev et al. 2009](#)). We thus link firm-level data on export deals to the procurement contracts awarded to these firms, and we test whether receipt of a politically motivated public

TABLE VIII
THE EFFECT OF A PUBLIC SECURITY AI CONTRACT ON AI EXPORTS

	Newly exporting firm			
	(1)	(2)	(3)	(4)
Public security	0.032* (0.018)	0.036** (0.017)	0.039** (0.019)	0.036** (0.018)
Contract quarter FE	No	Yes	Yes	Yes
Contract prefecture FE	No	Yes	Yes	Yes
Precontract software	No	No	Yes	Yes
Firm age	No	No	No	Yes

Notes. This table presents cross-sectional regressions at the firm level. The dependent variable is an indicator of whether the firm begins to export its AI products after receiving its first contract (i.e., the first difference in firm exporting status around the time of receiving a first government contract). The explanatory variable of interest is an indicator of whether the first contract was a (politically motivated) public security contract. The sample includes firms receiving their first contracts in prefectures that experienced above-median political unrest in the preceding quarter. Column (1) presents a simple regression; column (2) adds contract quarter and contract prefecture fixed effects; column (3) adds a control for firms' precontract software output; and column (4) adds a control for firms' age. Firms are weighted by their number of subsidiary firms. Robust standard errors are reported in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

security contract is associated with a change in a firm's status as an AI exporter. We compare the change in exporter status for firms receiving politically motivated public security contracts with firms receiving non-public security contracts in a politically sensitive environment to account for firm selection and the role of firms' political connections. Specifically, we examine the cross-sectional relationship between a change in exporter status around contract receipt, finding a significantly larger change among firms receiving a politically motivated public security contract. This is seen in the raw data (Table VIII, column (1)), and accounting for contract quarter fixed effects, contract prefecture fixed effects, firms' precontract software production, as well as firms' age (columns (2)–(4)). We observe a robust pattern that firms receiving politically motivated contracts are more likely (by 3.2 to 3.9 percentage points) to become exporters following receipt of these contracts. This effect is large: among AI firms that won at least one politically motivated public security contract, only 1% of them have exported their products between 2014 and 2021.

IV.C. Firm-Level Distortions Due to Politically Motivated Contracts

To the extent that politically motivated public security contracts may be accompanied by additional noncommercial demands

from the local government or may be associated with greater specialization, such contracts could differentially crowd out firms' commercial activities relative to the public security contracts that are not politically motivated but provide access to similar resources (e.g., data, capital, and political connections).⁵⁴ As discussed in [Beraja, Yang, and Yuchtman \(forthcoming\)](#), the greater the effects of politically motivated contracts on software production for the more general commercial market, the greater the effect these contracts would have on the trajectory of innovation in the AI sector.

To evaluate whether politically motivated contracts are associated with differential crowding out of commercial innovation, we compare the effects of politically motivated public security contracts to the effects of non-politically motivated public security contracts. We define politically motivated contracts as those issued following a quarter with above-median political unrest (as we did before), and politically neutral contracts as those issued following a quarter with below-median political unrest. We now limit our analysis only to public security contracts, and compare effects on software output of those plausibly granted out of political motivation with those that are more politically neutral.

[Online Appendix](#) Figure A.XV presents the coefficients indicating the differential effect of politically motivated public security contracts for the AI firms' commercial software production. We do not observe noticeable crowding out of commercial software production from politically motivated contracts. In fact, if anything, one sees that politically motivated contracts tend to induce firms to produce more commercial software, especially toward the later periods of the sampling frame.

IV.D. Cross-Firm Spillovers and Aggregate Effects

We consider the aggregate effects of politically motivated contracts. These may differ from the firm-level effects due to either positive or negative spillovers to other firms not receiving contracts. Positive spillovers across firms might arise because of knowledge spillovers from contracted firms to others (or across subsidiaries in the same mother firm). These spillovers may occur primarily in the locations where unrest occurs, or where contracted firms are headquartered and innovative activities may

54. This could arise from differential costs associated with developing products specifically for politically sensitive and demanding environments.

be concentrated. On the contrary, negative spillovers may arise if critical resources such as investments and human capital are disproportionately allocated to firms that have been awarded procurement contracts, or due to business-stealing effects among firms in the industry or across subsidiaries.⁵⁵

Gauging such spillovers along three margins, we examine AI innovation among firms never receiving procurement contracts that are: (i) headquartered in localities that have experienced political unrest; (ii) headquartered in localities where AI firms receiving politically motivated contracts are also headquartered; or (iii) part of a parent firm with other subsidiaries that have received politically motivated contracts.⁵⁶

Specifically, we estimate event study models in which the innovation of firms that never receive contracts is examined around (i) a quarter when the local prefecture experiences political unrest; (ii) a quarter when a politically motivated public security contract is issued to other firms headquartered in the same prefecture; and (iii) a quarter when a politically motivated public security contract is issued to another subsidiary of the same parent firm. We present the cumulative spillover effects two years after the relevant events in [Table IX](#), Panels A to C, respectively; we plot the full event study estimates from two years leading up to the event to two years after in [Online Appendix](#) Figure A.XVI. We find no evidence of negative spillovers in any of these cases. In fact, we find suggestive evidence of positive spillovers to the amount of commercial software produced by noncontracted firms.

Although these tests are not absolutely conclusive, the absence of evidence of significant distortions—at the firm level and across firms—as a result of autocrats' politically motivated procurement of AI technology suggests a positive aggregate effect on frontier AI innovation.⁵⁷

V. CONCLUDING THOUGHTS: THE IMPLICATIONS OF AI-TOCRACY

We document a mutually reinforcing relationship between facial-recognition AI innovation and China's autocratic regime.

55. In addition, firms not receiving contracts may be either positively or negatively affected by broader local policy changes enacted in response to local unrest.

56. These three margins are not intended to be an exhaustive catalog of spillovers, but are the important ones that we have the capacity to evaluate.

57. It is important to note that by examining only firms in facial-recognition AI, we are unable to investigate whether the increased frontier innovation in this field imposes costs on AI innovation in directions other than facial recognition or other fields beyond AI as a whole.

TABLE IX
 SPILLOVER EFFECTS OF POLITICALLY MOTIVATED PUBLIC SECURITY CONTRACTS ON
 SOFTWARE PRODUCTION

	Software produced		
	Total software (1)	Government (2)	Commercial (3)
Panel A: Firms headquartered in localities that experienced political unrest			
8 quarters after unrest	23.968** (10.122)	1.372 (1.102)	9.812** (4.709)
Panel B: Firms headquartered in localities where AI firms receiving politically motivated contracts are also headquartered			
8 quarters after contract	0.003 (0.071)	0.028 (0.033)	0.047 (0.046)
Panel C: Firms that are part of a mother firm with other subsidiaries that have received politically motivated contracts			
8 quarters after contract	-0.498 (0.381)	-0.102 (0.391)	1.365*** (0.310)

Notes. The table presents regression coefficients from an event study analysis of the software output of facial-recognition AI firms that do not receive politically motivated contracts. Panel A studies the effect of political unrest in the location where firms are headquartered. Panel B studies the effect of the receipt of a politically motivated contract by other firms on the firms not receiving contracts, but headquartered in the same prefecture. Panel C studies the effect of the receipt of a politically motivated contract by other firms on the firms not receiving contracts, but belonging to the same mother firm. All columns control for time period fixed effects and firm fixed effects. Panels B and C also control for contract fixed effects. Column (1) considers total software production as the outcome, column (2) considers government software, and column (3) considers commercial software. Standard errors are clustered at the prefecture level. * $p < .10$, ** $p < .05$, *** $p < .01$.

This relationship has direct implications for China's economic and political trajectories. First, China's autocratic politics may not constrain its ability to continue to push out the technological frontier in AI; rather, innovation in AI may be stimulated precisely because of China's autocratic politics. Second, continued frontier innovation and economic development in China may not be associated with more inclusive political institutions; rather, such innovation may further entrench the autocratic regime.

It is important to consider the extent to which our results generalize. While many technologies would not exhibit forces that generate mutually reinforcing relationships between autocracy and frontier innovation, the key forces that we highlight could shed light on prominent historical episodes of frontier innovation in, for example, the Soviet Union and imperial Germany. More generally, the evidence speaks to how state-sponsored innovation is supported in democracies, including innovation supported by DARPA in the United States, the high-tech sector supported by the military in Israel, and nuclear engineering programs led by the French state.

Looking ahead, a mutually reinforcing relationship between AI and autocracy may become relevant in other contexts. Russia, in particular, has already deployed facial-recognition AI for political control, and (not coincidentally) alongside China is among the world's leading producers of frontier facial-recognition AI technology.⁵⁸ Moreover, autocrats in other countries well inside the technological frontier may import Chinese AI technology for political control. Indeed, anecdotal evidence suggests that China's surveillance AI technology has already been exported to other autocracies.⁵⁹ One thus naturally worries that autocrat-supporting AI may beget more autocracies. The implications of China's AI innovation for the global political and economic landscape are worthy of further, rigorous investigation.

MIT AND NATIONAL BUREAU OF ECONOMIC RESEARCH, UNITED STATES

HARVARD UNIVERSITY, UNITED STATES

HARVARD UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH, UNITED STATES, AND CANADIAN INSTITUTE FOR ADVANCED RESEARCH, CANADA

LONDON SCHOOL OF ECONOMICS AND CENTRE FOR ECONOMIC POLICY RESEARCH, UNITED KINGDOM, AND CESIFO, GERMANY

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

58. [Online Appendix](#) Figure A.XVII presents the global ranking of the companies with the top 10 facial recognition algorithms in terms of prediction accuracy, as ranked by the Face Recognition Vendor Test (FRVT), organized by the National Institute of Standards and Technology (an agency of the U.S. Department of Commerce) and considered as one of the most authoritative AI industry competitions. Chinese firms occupy all of the top 4 positions; 5 out of the top 10 positions are occupied by Chinese and Russian firms. Regarding Russia's use of facial recognition for political control, see [Arkhipolov and Rudnitsky \(2021\)](#).

59. For example, according to an Atlantic article, "Xi Jinping is using artificial intelligence to enhance his government's totalitarian control—and he's exporting this technology to regimes around the globe. . . China is already developing powerful new surveillance tools, and exporting them to dozens of the world's actual and would-be autocracies." [Anderson \(2020\)](#).

DATA AVAILABILITY

The data underlying this article are available in the Harvard Dataverse, <https://doi.org/10.7910/DVN/GCOVGX> (Beraja et al. 2023).

REFERENCES

- Acemoglu, Daron, “Harms of AI,” NBER Working Paper 29247, 2021. <https://doi.org/10.3386/w29247>.
- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, “Artificial Intelligence and Jobs: Evidence from Online Vacancies,” *Journal of Labor Economics*, 40 (2022), S293–S340. <http://dx.doi.org/10.1086/718327>.
- Acemoglu, Daron, Simon Johnson, and James A. Robinson, “Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution,” *Quarterly Journal of Economics*, 117 (2002), 1231–1294. <http://dx.doi.org/10.1162/003355302320935025>.
- Acemoglu, Daron, and Pascual Restrepo, “Artificial Intelligence, Automation, and Work,” in *The Economics of Artificial Intelligence: An Agenda*, Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds. (Chicago: University of Chicago Press, 2019), 197–236. <http://dx.doi.org/10.7208/chicago/9780226613475.003.0008>.
- , “The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand,” *Cambridge Journal of Regions, Economy and Society*, 13 (2020), 25–35. <http://dx.doi.org/10.1093/cjres/rsz022>.
- Acemoglu, Daron, and James A. Robinson, “Economic Backwardness in Political Perspective,” *American Political Science Review*, 100 (2006), 115–131. <http://dx.doi.org/10.1017/S0003055406062046>.
- , *Why Nations Fail: The Origins of Power, Prosperity, and Poverty* (New York: Crown Business, 2012).
- , *The Narrow Corridor: States, Societies, and the Fate of Liberty* (New York: Penguin Books, 2020).
- Aghion, Philippe, Alberto Alesina, and Francesco Trebbi, “Democracy, Technology, and Growth,” in *Institutions and Economic Performance*, Elhanan Helpman, ed., (Cambridge, MA: Harvard University Press, 2008), 511–543. <http://dx.doi.org/10.4159/9780674038578-014>
- Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones, “Artificial Intelligence and Economic Growth,” in *The Economics of Artificial Intelligence: An Agenda*, Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds. (Chicago: University of Chicago Press, 2019), 237–290. <https://doi.org/10.7208/chicago/9780226613475.003.0009>.
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb, “Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction,” *Journal of Economic Perspectives*, 33 (2019), 31–50. <http://dx.doi.org/10.1257/jep.33.2.31>.
- Anderson, Michael L., “Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects,” *Journal of the American Statistical Association*, 103 (2008), 1481–1495. <http://dx.doi.org/10.1198/016214508000000841>.
- Andersen, Ross, “The Panopticon Is Already Here,” *The Atlantic* 2020, <https://www.theatlantic.com/magazine/archive/2020/09/china-ai-surveillance/614197/>.
- Arkipov, Ilya, and Jake Rudnitsky, “In Moscow, Big Brother Is Watching and Recognizing Protesters,” *Bloomberg*, 2021, https://www.bloomberg.com/news/articles/2021-05-02/in-moscow-big-brother-is-watching-and-recognizing-protesters_xj4y7vzkg.
- Bai, Chong-en, Chang-Tai Hsieh, and Zheng Song, “Special Deals with Chinese Characteristics,” *NBER Macroeconomics Annual*, 34 (2020), 341–379. <http://dx.doi.org/10.1086/707189>.

- Barro, Robert J., "Democracy and Growth," *Journal of Economic Growth*, 1 (1996), 1–27. <http://dx.doi.org/10.1007/BF00163340>.
- Beraja, Martin, Andrew Kao, David Y. Yang, and Noam Yuchtman, "Replication Data for: 'AI-tocracy'," (2023), Harvard Dataverse. <https://doi.org/10.7910/DVN/GCOVGX>.
- Beraja, Martin, David Y. Yang, and Noam Yuchtman, "Data-Intensive Innovation and the State: Evidence from AI firms in China," *Review of Economic Studies* (forthcoming). <http://dx.doi.org/10.1093/restud/rdac056>.
- Beraja, Martin, and Nathan Zorzi, "Inefficient Automation," NBER Working Paper 30154, 2022. <https://doi.org/10.3386/w30154>.
- Besley, Timothy, and Torsten Persson, "The Origins of State Capacity: Property Rights, Taxation, and Politics," *American Economic Review*, 99 (2009), 1218–1244. <http://dx.doi.org/10.1257/aer.99.4.1218>.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, 110 (2020), 1104–1144. <http://dx.doi.org/10.1257/aer.20180338>.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event Study Designs: Robust and Efficient Estimation," Working Paper, 2017. <https://doi.org/10.48550/arXiv.2108.12419>.
- Brandt, Loren, and Thomas G. Rawski, *China's Great Economic Transformation* (Cambridge: Cambridge University Press, 2008). <http://dx.doi.org/10.1017/CBO9780511754234.002>.
- Bursztyjn, Leonardo, Davide Cantoni, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang, "Persistent Political Engagement: Social Interactions and the Dynamics of Protest Movements," *American Economic Review: Insights*, 3 (2021), 233–250. <http://dx.doi.org/10.1257/aeri.20200261>.
- Cagé, Julia, and Valeria Rueda, "The Long-Term Effects of the Printing Press in Sub-Saharan Africa," *American Economic Journal: Applied Economics*, 8 (2016), 69–99. <http://dx.doi.org/10.1257/app.20140379>.
- Campante, Filipe R., Davin Chor, and Bingjing Li, "The Political Economy Consequences of China's Export Slowdown," *Journal of the European Economic Association*, (forthcoming). <http://dx.doi.org/10.1093/jeea/jvad007>.
- Campante, Filipe R., Ruben Durante, and Francesco Sobbrío, "Politics 2.0: The Multifaceted Effect of Broadband Internet on Political Participation," *Journal of the European Economic Association*, 16 (2018), 1094–1136. <http://dx.doi.org/10.1093/jeea/jvx044>.
- Cantoni, Davide, Yuyu Chen, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang, "Curriculum and Ideology," *Journal of Political Economy*, 125 (2017), 338–392. <https://doi.org/10.1086/690951>.
- Chen, Jidong, Jennifer Pan, and Yiqing Xu, "Sources of Authoritarian Responsiveness: A Field Experiment in China," *American Journal of Political Science*, 60 (2016), 383–400. <http://dx.doi.org/10.1111/ajps.12207>.
- Chen, Yuyu, and David Y. Yang, "The Impact of Media Censorship: 1984 or Brave New World?" *American Economic Review*, 109 (2019), 2294–2332. <http://dx.doi.org/10.1257/aer.20171765>.
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Dufo, Christian Hansen, Whitney Newey, and James Robins, "Double/Debiased Machine Learning for Treatment and Structural Parameters," *Econometrics Journal*, 21 (2018), C1–C68. <http://dx.doi.org/10.1111/ectj.12097>.
- Chin, Josh, and Liza Lin, *Surveillance State: China's Quest to Launch a New Era of Social Control* (New York: St. Martin's Press, 2022).
- Cowgill, Bo, and Catherine Tucker, "Algorithmic Fairness and Economics," Columbia Business School Research Paper, 2020.
- Dell, Melissa, "The Persistent Effects of Peru's Mining Mita," *Econometrica*, 78 (2010), 1863–1903. <https://doi.org/10.3982/ECTA8121>.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, "What Do We Learn from the Weather? The New Climate-Economy Literature," *Journal of Economic Literature*, 52 (2014), 740–798. <http://dx.doi.org/10.1257/jel.52.3.740>.

- Dell, Melissa, and Benjamin A. Olken, "The Development Effects of the Extractive Colonial Economy: The Dutch Cultivation System in Java," *Review of Economic Studies*, 87 (2020), 164–203. <http://dx.doi.org/10.1093/restud/rdz017>.
- Donaldson, Dave, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, 108 (2018), 899–934. <http://dx.doi.org/10.1257/aer.20101199>.
- Eeckhout, Jan, and Laura Veldkamp, "Data and Market Power," NBER Working Paper 30022, 2022. <https://doi.org/10.3386/w30022>.
- Enikolopov, Ruben, Alexey Makarin, and Maria Petrova, "Social Media and Protest Participation: Evidence from Russia," *Econometrica*, 88 (2020), 1479–1514. <http://dx.doi.org/10.3982/ECTA14281>.
- Farboodi, Maryam, and Laura Veldkamp, "A Model of the Data Economy," NBER Working Paper 28427, 2022. <https://doi.org/10.3386/w28427>.
- Feldstein, Steven, *The Global Expansion of AI Surveillance* (Washington, DC: Carnegie Endowment for International Peace, 2019).
- Filatotchev, Igor, Xiaohui Liu, Trevor Buck, and Mike Wright, "The Export Orientation and Export Performance of High-Technology SMEs in Emerging Markets: The Effects of Knowledge Transfer by Returnee Entrepreneurs," *Journal of International Business Studies*, 40 (2009), 1005–1021. <http://dx.doi.org/10.1057/jibs.2008.105>.
- Glaeser, Edward L., Giacomo A. M. Ponzetto, and Andrei Shleifer, "Why Does Democracy Need Education?," *Journal of Economic Growth*, 12 (2007), 77–99. <http://dx.doi.org/10.1007/s10887-007-9015-1>.
- Goldfarb, Avi, and Daniel Trefer, "AI and International Trade," NBER Working Paper 24254, 2018. <https://doi.org/10.3386/w24254>.
- Greitens, Sheena Chestnut, "Dealing with Demand for China's Global Surveillance Exports," Brookings Institution Global China Report, 2020.
- Guriev, Sergei, Nikita Melnikov, and Ekaterina Zhuravskaya, "3G Internet and Confidence in Government," *Quarterly Journal of Economics*, 136 (2021), 2533–2613. <http://dx.doi.org/10.1093/qje/qjaa040>.
- Guriev, Sergei, and Daniel Treisman, "Informational Autocrats," *Journal of Economic Perspectives*, 33 (2019), 100–127. <http://dx.doi.org/10.1257/jep.33.4.100>.
- Huillery, Elise, "History Matters: The Long-Term Impact of Colonial Public Investments in French West Africa," *American Economic Journal: Applied Economics*, 1 (2009), 176–215. <http://dx.doi.org/10.1257/app.1.2.176>.
- Jones, Charles I., and Christopher Tonetti, "Nonrivalry and the Economics of Data," *American Economic Review*, 110 (2020), 2819–2858. <http://dx.doi.org/10.1257/aer.20191330>.
- King, Gary, Jennifer Pan, and Margaret E. Roberts, "How Censorship in China Allows Government Criticism but Silences Collective Expression," *American Political Science Review*, 107 (2013), 326–343. <http://dx.doi.org/10.1017/S0003055413000014>.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan, "Human Decisions and Machine Predictions," *Quarterly Journal of Economics*, 133 (2018), 237–293. <http://dx.doi.org/10.1093/qje/qjx032>.
- Korinek, Anton, and Joseph E. Stiglitz, "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," in *The Economics of Artificial Intelligence: An Agenda*, Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds. (Chicago: University of Chicago Press, 2018), 349–390. <http://dx.doi.org/10.7208/chicago/9780226613475.003.0014>.
- Lau, Lawrence J., Yingyi Qian, and Gerard Roland, "Reform without Losers: An Interpretation of China's Dual-Track Approach to Transition," *Journal of Political Economy*, 108 (2000), 120–143. <http://dx.doi.org/10.1086/262113>.
- Lipset, Seymour Martin, "Some Social Requisites of Democracy: Economic Development and Political Legitimacy," *American Political Science Review*, 53 (1959), 69–105. <http://dx.doi.org/10.2307/1951731>.
- Lorentzen, Peter, "China's Strategic Censorship," *American Journal of Political Science*, 58 (2014), 402–414. <http://dx.doi.org/10.1111/ajps.12065>.

- Lowes, Sara, and Eduardo Montero, "Concessions, Violence, and Indirect Rule: Evidence from the Congo Free State," *Quarterly Journal of Economics*, 136 (2021), 2047–2091. <http://dx.doi.org/10.1093/qje/qjab021>.
- Madestam, Andreas, Daniel Shoag, Stan Veuger, and David Yanagizawa-Drott, "Do Political Protests Matter? Evidence from the Tea Party Movement," *Quarterly Journal of Economics*, 128 (2013), 1633–1685. <http://dx.doi.org/10.1093/qje/qjt021>.
- Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71 (2003), 1695–1725. <http://dx.doi.org/10.1111/1468-0262.00467>.
- Mokyr, Joel, "Punctuated Equilibria and Technological Progress," *American Economic Review*, 80 (1990), 350–354. <https://www.jstor.org/stable/2006599>.
- Moretti, Enrico, Claudia Steinwender, and John Van Reenen, "The Intellectual Spoils of War? Defense R&D, Productivity and International Spillovers," *Review of Economics and Statistics*, forthcoming. <http://dx.doi.org/10.1162/rest.a.01293>.
- Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny, "Industrialization and the Big Push," *Journal of Political Economy*, 97 (1989), 1–25. <http://dx.doi.org/10.1086/261641>.
- North, Douglass C., John Joseph Wallis, and Barry R. Weingast, *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History* (Cambridge: Cambridge University Press, 2009). <http://dx.doi.org/10.1017/CBO9780511575839>.
- North, Douglass C., and Barry R. Weingast, "Constitutions and Commitment: The Evolution of Institutions Governing Public Choice in Seventeenth-Century England," *Journal of Economic History*, 49 (1989), 803–832. <https://doi.org/10.1017/S0022050700009451>.
- Nunn, Nathan, "The Long-Term Effects of Africa's Slave Trades," *Quarterly Journal of Economics*, 123 (2008), 139–176. <http://dx.doi.org/10.1162/qjec.2008.123.1.139>.
- Olson, Mancur, "Dictatorship, Democracy, and Development," *American Political Science Review*, 87 (1993), 567–576. <http://dx.doi.org/10.2307/2938736>.
- Qin, Bei, David Strömberg, and Yanhui Wu, "Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda," *Journal of Economic Perspectives*, 31 (2017), 117–140. <http://dx.doi.org/10.1257/jep.31.1.117>.
- , "Social Media and Protests in China," CEPR Discussion Paper No. DP16731, 2020.
- Roberts, Margaret E., *Censored: Distraction and Diversion Inside China's Great Firewall* (Princeton, NJ: Princeton University Press, 2018). <http://dx.doi.org/10.23943/9781400890057>.
- Schwab, Klaus, *The Fourth Industrial Revolution* (Redfern, Sydney: Currency, 2017).
- Shirk, Susan L., *China: The Fragile Superpower* (Oxford: Oxford University Press, 2007).
- Shleifer, Andrei, and Robert W. Vishny, *The Grabbing Hand: Government Pathologies and Their Cures* (Cambridge, MA: Harvard University Press, 2002).
- Song, Zheng, Kjetil Storesletten, and Fabrizio Zilibotti, "Growing Like China," *American Economic Review*, 101 (2011), 196–233. <http://dx.doi.org/10.1257/aer.101.1.196>.
- Tirole, Jean, "Digital Dystopia," *American Economic Review*, 111 (2021), 2007–2048. <http://dx.doi.org/10.1257/aer.20201214>.
- Treisman, Daniel, "Economic Development and Democracy: Predispositions and Triggers," *Annual Review of Political Science*, 23 (2020), 241–257. <http://dx.doi.org/10.1146/annurev-polisci-050718-043546>.
- Tsai, Lily L., *Accountability without Democracy: Solidary Groups and Public Goods Provision in Rural China* (Cambridge: Cambridge University Press, 2007). <http://dx.doi.org/10.1017/CBO9780511800115>.

- Valencia Caicedo, Felipe “The Mission: Human Capital Transmission, Economic Persistence, and Culture in South America,” *Quarterly Journal of Economics*, 134 (2019), 507–556. <http://dx.doi.org/10.1093/qje/qjy024>.
- Vernon, Raymond, “International Investment and International Trade in the Product Cycle,” *Quarterly Journal of Economics*, 80 (1966), 190–207. <http://dx.doi.org/10.2307/1880689>.
- Wang, Zhi, Qinghua Zhang, and Li-An Zhou, “Career Incentives of City Leaders and Urban Spatial Expansion in China,” *Review of Economics and Statistics*, 102 (2020), 897–911. <http://dx.doi.org/10.1162/rest.a.00862>.
- Yang, Jing, and Lingling Wei, “China’s President Xi Jinping Personally Scuttled Jack Ma’s Ant IPO,” *Wall Street Journal*, 2020, <https://www.wsj.com/articles/china-president-xi-jinping-halted-jack-ma-ant-ipo-11605203556>.
- Zhang, Daniel, Saurabh Mishra, Erik Brynjolfsson, John Etchemendy, Deep Ganguli, Barbara Grosz, Terah Lyons, James Manyika, Juan Carlos Niebles, Michael Sellitto, Yoav Shoham, Jack Clark, and Raymond Perreault, “The AI Index 2021 Annual Report,” AI Index Steering Committee, Human-Centered AI Institute, Stanford University Technical Report, 2021. <https://doi.org/10.48550/arXiv.2103.06312>.
- Zuboff, Shoshana, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: PublicAffairs, 2019).