Research paper

Juror interpretations of metadata and content information: implications for the going dark debate

Anne E. Boustead [1,*] and Matthew B. Kugler[2]

[1]School of Government & Public Policy, University of Arizona, Tucson, AZ 85719, USA and [2]Northwestern Pritzker School of Law, Northwestern University, Chicago, IL 60611, USA

*Correspondence address. School of Government & Public Policy, University of Arizona, 1145 South Campus Drive, #331, Tucson, AZ 85719, USA. E-mail: boustead@arizona.edu

Abstract

The rise of consumer encryption has led to a fierce debate over whether the loss of potential evidence due to encryption will be offset by the increase in evidence available from electronic metadata. One major question raised by this debate is how jurors will interpret and value metadata as opposed to content information. Though there are plausible arguments in favor of the persuasive power of each type of evidence, to date no empirical study has examined how ordinary people, potential jurors, view each of these sorts of evidence.
We address this issue through a series of survey experiments that present respondents with hypothetical criminal trials, randomly assigning them to descriptions featuring either metadata or content information. These studies show that the relative power of content and metadata information is highly contextual. Content information and metadata can be equally useful when conveying logically equivalent information. However, content information may be more persuasive where the defendant's state of mind is critical, while metadata can more convincingly establish a pattern of behavior. This suggests that the rise of encryption will have a heterogeneous effect on criminal cases, with the direction of the effect depending on the facts that the prosecution must prove.

Key words: surveillance, technology, criminal investigations, jurors, decision-making

Introduction

There is a contentious debate over whether developers of encrypted communication products should be required to ensure that government officials are able to access the contents of these communications. On one hand, advocates for government access cite the critical role that communication contents play in criminal and national security investigations and the substantial public safety harms that may occur if these investigations are slowed or stopped [1]. On the other hand, advocates for strong consumer encryption discuss the myriad ways this technology can contribute to a safe and fair society—from promoting cybersecurity [2] to safeguarding the activities of minoritized groups [3]. Advocates for strong consumer encryption frequently also cite the availability of other sources of information—including cell phone metadata and data from Internet of Things (IoT)

devices—as potentially mitigating the impediments to investigations caused by consumer-grade encryption.

A significant portion of the policy debate over whether to ensure government access to encrypted devices has focused on the availability of substitutes—other forms of surveillance or sources of information that can be used in place of the information lost due to encryption. The effectiveness of potential substitutes for now-encrypted communications largely depends on how actors in the criminal justice system behave when encountering the contents of communications in comparison to their behavior when encountering potential substitutes. To date, however, few studies have attempted to empirically describe whether and how actors in the criminal justice system behave differently when encountering the contents of information versus potential substitutes.

In this article, we attempt to address this question in the context of one particularly important set of criminal justice decision-makers: the jurors who are ultimately asked to render verdicts in criminal cases. While some psychological theories of juror decision-making would seem to suggest that content information could be particularly convincing, other research suggests that jurors might be predisposed to trust or even expect technically derived evidence such as metadata. To test these competing hypotheses, we use a series of survey experiments to provide the first empirical evidence regarding whether respondents acting as jurors in a hypothetical criminal case reach different conclusions when presented with content information or comparable metadata. We find that whether respondents are more likely to indicate they would convict when presented with content information or metadata is highly context-dependent. Content information and metadata appear to be equally convincing when they present logically equivalent information. But content information appears to be more convincing when establishing specifics about the defendant's state of mind is critical to the prosecution's case, and metadata appears to be more convincing when the prosecution's case hinges on establishing a pattern of behavior. Importantly, these findings suggest that a shift from presenting content information to metadata at trial may have a differential effect on juror decision-making across different types of crimes.

The remainder of this article proceeds in four parts. Part 1 reviews the current literature on the information available during criminal investigations and trials that are likely to be impacted by two broad technological trends: the inability of government actors to access the contents of information due to encryption (sometimes referred to as "going dark") and a dramatic increase in the scope and types of information routinely collected due to commercial actors (sometimes referred to as "the golden age of surveillance"). We then explore how the uncertainty about these trends is likely to impact juror decision-making, concluding that current theoretical models and empirical evidence do not provide a clear answer to this question. Part 2 describes our experimental methods and results. We then characterize the overall findings and implications of our studies in Part 3 and conclude with a few specific policy recommendations in Part 4.

Technological trends will impact the information available for criminal investigations and trials

Today, the sources of information available during criminal investigations and trials are being reshaped by two competing technological trends [4]. On one hand, the increased use of encryption on consumer devices means that common sources of content information may be becoming inaccessible to law enforcement, or "going dark." On the other hand, new forms of consumer technology may make detailed, cheaply accessible information available during criminal investigations to a previously unimaginable extent, leading to a "golden age" of surveillance. In this part, we explore the arguments surrounding both the "going dark" and "golden age" perspectives on recent trends in surveillance. We then turn to the potential intersection between these trends, investigating whether information derived through metadata and other emerging sources of information can compensate for the information lost due to encrypted communications. Finally, we conclude by discussing the potential impact of these trends on juror decision-making.

Going dark

While individuals have long used codes to protect the contents of their communications [5], encryption required significant sophisti-

cation to implement and consequently was not widely used [6]. In recent years, however, changes in consumer technology have made it dramatically easier for individuals to encrypt their communications [7]. Apple enabled encryption by default with iOS 8 in 2014, and strengthened their encryption implementation to make it "not technically feasible…to respond to government warrants for the extraction of this data from devices in their possession" [8]. Google soon adopted a similar practice for their Android operating system [9]. Use of encryption consequently increased in both the smartphone context and in general [10]. The resulting spread of encryption has made it much more likely that law enforcement will be unable to access the contents of a device they have obtained during the course of a criminal investigation, even with a warrant [11], leading to a phenomenon referred to as "going dark" [12].

According to analysts, advocates, and representatives from the law-enforcement community, the loss of content information due to encryption has had a substantial and detrimental impact on criminal investigations. Former US Attorney General William Barr has said, "the costs of irresponsible encryption that blocks legitimate law enforcement access is ultimately measured in a mounting number of victims…[and] crimes that could have been prevented if law enforcement had been given lawful access to encrypted evidence" [13]. Because of barriers created by encryption technologies, federal law enforcement has stated it "may not be able to identify and stop terrorists who are using social media to recruit, plan, and execute an attack" [14].

Though much of the conversation around the costs of going dark focuses on implications for national security and terrorism investigations, this is by no means the only type of investigation impacted by the loss of evidence due to encryption. Prosecutions for possession of child pornography, e.g , become very difficult when law enforcement cannot access the contents of a device to show they include child pornography [1], and drug trafficking investigations are hindered when cartels utilize encrypted communication mechanisms to organize narcotics shipments [13].

Despite the political salience of arguments that law enforcement is losing access to crucial information during criminal investigations due to encryption, there is scant empirical evidence available on the scope or impact of this phenomenon. For example, the Wiretap Reports present yearly statistics on the instances where law enforcement encounters encryption when attempting to execute a Title III warrant [15]. Similarly, the New York District Attorney's Office has released a series of reports on encounters with encryption in their jurisdiction, although the measures reported vary from year to year in ways that may make comparisons difficult [16]. But these data in one case consider only a single form of surveillance and in the other only a single jurisdiction, making generalization difficult. Furthermore, these data may not show the full impact of encryption on law enforcement access to data, as law enforcement may preemptively decide not to undertake the costs associated with conducting surveillance in contexts where they believe encryption will make it unlikely they will obtain information [17].

Golden age of surveillance

In contrast to the "going dark" narrative, other technological trends are creating new options for surveillance. In the eyes of some, we now live in a "golden age of surveillance," where "investigatory agencies have unprecedented access to information about a suspect" [18]. Though this "golden age" narrative does not assert that criminal investigations will be unimpeded by the loss of content information to encryption, it argues that a number of technological trends have dra-

matically increased the scope of information available during criminal investigations and prosecutions [19].

One of these trends is the rise in new consumer technologies—notably IoT devices—that generate new sources of information and data records that may be useful in criminal investigations. The IoT is comprised of myriad "smart" physical devices [20], potentially capable of both conveying and receiving information through the Internet [21]. IoT devices routinely collect information about nearby activities, sometimes without active participation, awareness, or control from the individuals involved [22]. IoT devices frequently contain an array of sensors capable of recording information about their environment, providing both rich sources of already collected information and opportunities for law enforcement to conduct active surveillance using sensors already placed in their target's environment [19]—and raising privacy concerns [23]. According to one former prosecutor, reimagining how technological shifts might have changed an investigation she participated in during the 1990s, "it would have been worth considering how IoT devices could have been exploited in [the suspects'] apartments or the cars they drove," which could have allowed the recording of communications without the risks associated with obtaining physical access to a space in order to plant a bug [24]. One can also use the location information generated by cellphones in place of installing tracking devices on people and vehicles.

For the purposes of this paper, we refer to the type of information generated by these IoT devices and, particularly, by cellphones, as metadata. This terminology is inexact. Metadata are merely data that describe other data [25]. For example, metadata about a Microsoft Word document might include the date it was created, the size of the file, and the last date it was accessed, while content information would include the substance of the document itself. This distinction between metadata and content information may be less clear in more complex operations [26]. In practical terms, however, there is a relatively clear distinction between the kinds of information that may be increasingly encrypted and therefore unavailable to law enforcement—generally message and call content—and the types of information that are now automatically collected by third-party providers and stored in their records—generally some form of metadata.

As consumer data have become a valuable commodity, companies are now strongly incentivized to ensure that the data they collect about customers can be easily accessed and used for authorized purposes [27]. These incentives may ensure that sources of individual-level data will continue to be available for use in criminal investigations and trials, even if communication content information becomes unavailable due to encryption [19]. This seems especially likely in contexts where companies are able to generate revenue through providing police with access to information [28]. For example, Clearview AI used data from a range of publicly available sources (including Venmo and Facebook) to develop a facial recognition tool that has been used by more than a thousand law-enforcement agencies in the USA [29].

Though some companies are embracing law-enforcement compatibility, others are resisting the notion that the collection of detailed, unencrypted, individual-level data is a necessary component of commercializing these data. For example, Apple recently updated their iOS operating system to require app developers to explicitly and obviously seek permission prior to tracking user data [30], which may reduce the amount of individual-level data collected and therefore available to law enforcement. In addition, there has recently been increasing recognition of the role that large technology companies may place as surveillance intermediaries: companies with "the incentives and means to meaningfully constrain government surveillance" [31].

Potential impact of the "golden age" on "going dark"

There has been an ongoing debate over the extent to which evidence from other sources can serve as at least a partial substitute for communications made inaccessible to law enforcement by encryption. The extent to which actors may substitute one activity for another, or consumption of one good for another, is a commonly considered question in economics, law, and public policy [32–34]. Influenced by this literature, we are primarily concerned with a behavioral understanding of substitutes.

If so, then law enforcement's increased ability to access other sources of information may ease the challenges posed by traditional sources of content information going dark—with potential implications for the debate over policy responses to this phenomenon. From a technological perspective, the arguments for replacement of communications that cannot be obtained by law enforcement due to encryption with alternative sources of information are straightforward: Other forms of surveillance may convey data that are similar to the information content lost due to encryption. For example, even if law enforcement can no longer access communications due to the use of encrypted devices, they may be able to utilize IoT devices to record similar conversations through a different mechanism. Although the two conversations may not be identical, the "inability to monitor an encrypted channel could be mitigated by the ability to monitor from afar a person through a different channel" [19].

But both representatives from the law-enforcement community and some scholars have argued that this substitution potential has been overstated [35]. Though, metadata are "especially valuable for providing information about 'who,' 'where,' and 'when,'…[w]ithout access to the contents of messages and stored data, one cannot directly determine 'what'—such as the plans and intentions of criminals" [36]. For example, it is less likely to reveal information about the motives underlying criminal activity [24]. While Cloud backups may be useful for obtaining unencrypted versions of data stored on encrypted devices, backups cannot be uploaded to the Cloud if the device is not connected to the Internet—potentially limiting their utility to law enforcement [37].

Additionally, arguments about how technically commensurate two sources of information are frequently rest on (stated or unstated) assumptions about the behaviors of various groups of criminal justice actors who interact with the information in question. Importantly, variation in behavior within these groups may mean that other forms of surveillance may act as a substitute for information lost to encryption only under certain conditions or for particular subgroups. For example, while criminals may elect to discontinue Cloud backups of incriminating communications in order to ensure they are not available in unencrypted formats, this planning is only possible when the perpetrator is planning their criminal activities in advance and recognizes that their communications may be incriminating [4]. Similarly, while law enforcement's ability to use another form of surveillance in place of the contents of communications lost due to encryption is dependent on them having knowledge that the alternative form of surveillance exists and being able to obtain information through it, this process may be easier for larger departments with more specialized officers—allowing only some agencies to take advantage of the golden age of surveillance [38]. Consequently, it is difficult to determine whether and under what circumstances criminal justice actors would be able to use metadata in place of the contents of encrypted communications.

These technological trends will have an uncertain impact on juror decision-making

One major outstanding question in this debate is how jurors, who may ultimately be presented with metadata evidence in the place of content evidence, will respond to this shift. As juror decision-making depends on the evidence presented to them, technological changes that shift the sources of evidence available for use in criminal trials may in turn change the conclusions reached by jurors in these cases—even when the sources of evidence available convey similar information. However, it is not clear in advance what direction this turn will take.

Content information may be particularly easy for jurors to understand. Communication content is a regular feature of day-to-day life, making it easier for jurors to incorporate this evidence into their understanding of the case as a whole. The psychological literature frequently describes juror decision-making through a "Story Model," where jurors use the evidence presented at trials to construct what they believe to be the most credible narrative that describes what happened, process the legal information given through jury instructions as a series of decision alternatives, and then select a verdict by deciding which options best fit the narrative of the case [39]. Under this model, content information may be inherently more persuasive than metadata because jurors find it easier to construct a narrative around a conversation than a technical abstraction [40]. Stories about human behavior frequently and prominently feature communications, and individuals frequently communicate with others in the course of their day-to-day lives. When presented with contents of communications, jurors may both intuitively understand and value the evidence.

The Story Model also suggests that jurors select the most credible narrative by the extent to which each narrative fits with the evidence presented, lack of internal contradictions, completeness, plausibility, and uniqueness [41]. These criteria may be more easily satisfied by content information than metadata, even when the two convey similar information. In particular, content information may provide a richer and more detailed picture of interactions between two individuals, reducing ambiguity and facilitating juror efforts to incorporate this evidence into their understanding of the case. The empirical literature suggests that content information will be more effective at achieving these aims when it presents coherent information [42]. In contrast, metadata may be considered inherently less complete—and therefore less convincing—as it does not include information about the substance of the transaction it describes [39].

Consider a hypothetical criminal case in which the jury is presented with a transcript of text messages showing that the victim recently paid a sum of money to the defendant. Jurors may not only evaluate the information conveyed by the text messages but also consider the tone of the messages. Even when the text messages convey bare-bones information that would logically be identical to metadata indicating that a payment had been made between the two ("I just paid you $500"), the fact that the victim felt the need to convey that in a text message may have implications for the relationship between the defendant and the victim. From a narrative perspective, content information may provide much richer information that jurors can use to build out narratives of the case and evaluate their plausibility.

Additionally, content information may be used to show not only that the defendant committed the crime in question but also their mental state while doing so. Though metadata may be useful for showing that the defendant was at the scene of the crime, content information may also show the defendant's awareness of the wrong-fulness of their actions, their emotions about the crime, and whether they demonstrated remorse prior to being held to account at trial—all of which are relevant for understanding the blameworthiness of the defendant's actions and determining their punishment. The potential richness of content data suggests that jurors may, on average, recommend longer sentences when exposed to content information as opposed to metadata, particularly when the content information speaks to the defendant's emotional state or awareness of the wrong-fulness of their actions.

On the other hand, jurors may have come to expect technological evidence like metadata, with its appearance of sophistication and reliability, and therefore find it especially persuasive. There has long been discussion over whether viewing popular media portrayals of crime scene investigators may cause jurors to expect extensive forensic evidence to be presented at trial [43]. Often coined the "CSI Effect," there is limited and mixed empirical evidence that such a phenomenon exists in practice [44–46]. Similarly, it is possible that widespread news reports about government surveillance programs could create expectations that extensive evidence from surveillance would be available and could be presented at trial. "[A]s jurors come to recognize how much data is (theoretically) available to law enforcement, they will view prosecutorial cases more skeptically if they're not presented with all that data" [35]. Under these arguments, jurors may therefore to come to expect that information derived from metadata be presented at trial—especially after the Snowden disclosures brought attention to the potential scope of metadata collection [47]. Metadata may prove sufficiently convincing to make up for the reduced availability of content information at criminal trials due to encryption, or may even prove more convincing than content information under certain circumstances.

There have been few empirical studies that consider behavioral responses to encrypted devices by criminal justice actors, and we could identify none that considered how the widespread use of consumer encryption might impact juror decision-making. For example, Pell [24] provides an interesting case study of how a case investigated and prosecuted in 1999 might be impacted by both the rise of encrypted communications and the information newly available to law enforcement due to the IoT, including consideration of how investigators may have behaved under both circumstances. However, juror decision-making has been broadly studied from a variety of theoretical [48], empirical [49], and policy [50] perspectives, resulting in a well-developed set of methodologies for answering questions about these behaviors. Consequently, the time is ripe for an empirical investigation of the comparative effects of content information and metadata on jury decision-making.

Experimental methods and results

To explore the ways in which potential jurors would understand different forms of evidence, we conducted a series of vignette-based studies. Each study described a criminal trial. Participants were given a description of the main evidence, asked to read a set of jury instructions, and then completed a verdict sheet. They also made a variety of ratings to further explain their verdict. For the screens presenting the major evidence, participants were required to spend at least 30 seconds on the page. They could, however, spend as long as they liked and even return to the page if they so desired.

Participants for each study were recruited via CloudResearch, an Internet panel management company [51]. This allowed for a cheap if not entirely representative sample. Quotas were set for each study

to require an approximately even split between men and women and to fix a median age in the 40s, consistent with the national median for adults. Exact splits varied from study to study due to excess responses provided by the sample provider. To complete the study, participants had to respond appropriately to an attention check item that requested a particular response and complete a CAPTCHA. Data from a small number of participants were also discarded if they finished the study in less than one-third of the median completion time or if they wrote gibberish in response to the question asking them to explain their verdict (0.7% in Study 1; 2.2% in Study 2; and 2.6% in Study 3).

Experiment series 1: form of evidence does not always matter

For the first study, the primary goal was to determine whether participants would understand metadata and content data as conveying substantively the same information, with the same level of reliability, when logically they should. This tests whether there is a bias for, or against, either type of data.

A total of 270 participants were recruited via CloudResearch and passed attention checks. They were 133 men, 136 women, and 1 person reporting another gender identification. They had a median age of 45 and 50.7% had college degrees, making them somewhat more educated than national norms.

The scenarios these participants read described an armed robbery of a convenience store. The defendant at trial matched the general build of the perpetrator and owned a gun of the same model, but no witness could positively identify him and no fingerprint or DNA evidence was presented. The defendant was also said to have a criminal record of theft from several years prior.

Further evidence in the scenario varied by condition. First, either metadata or content data placed the defendant at the scene of the crime. In the metadata case, an officer testified that the defendant's cellphone signal showed them to be within 1 block of the relevant store within minutes of the robbery. In the content condition, an officer testified that WhatsApp messages between the defendant and the defendant's girlfriend had been recovered. These messages included the defendant telling his girlfriend that he had "just been at the store on [street of crime]" shortly after the crime occurred.

The second factor that varied across conditions was the number of robberies at issue. The defendant was being charged with either one or five counts of robbery with a firearm. In the condition where five counts were charged, it was said that the evidence was fundamentally the same for each robbery. Either metadata placed the defendant at the scene of the robbery within minutes of the robbery each time or the defendant sent a similar message to his girlfriend each time. This was therefore a 2 × 2 design. The full text of the single count metadata condition is included in the Online Appendix.

A series of factorial ANOVAs were conducted on the primary response variables. On a percentile scale (0–100), participants reported that the defendant was much more likely to be guilty when he had been at the scene of five robberies ($M = 64.94$, SD = 27.77) than when he had been at the scene of only one ($M = 48.47$, SD = 29.22), $F(1, 217) = 18.13$, $P < 0.001$, $\eta^2 = 0.08$. Whether the evidence was metadata or content data had no effect, $F(1, 217) = 0.00$, $\eta^2 = 0.00$, and there was no interaction, $F(1, 217) = 0.18$, $\eta^2 = 0.00$. This translated into a conviction rate of 29.3% in the one-count condition and 49.5% in the five-count condition.

Participants were also asked to rate how much weight they put on each of four types of evidence: physical similarities between the defendant and perpetrator, similarities between the gun used by the perpetrator and that owned by the defendant, the defendant's prior convictions, and the cellphone evidence (varying by condition) on scales ranging from 0—no weight to 10—a great deal of weight. There were no significant effects of cellphone evidence type (meta or content data) on the weight assigned to any piece of evidence. For the weight assigned to cellphone evidence, the difference between conditions did not even approach significance. $F(1, 217) = 1.20$, $P = 0.27$, $\eta^2 = 0.01$ (Meta $M = 5.95$, SD = 2.97; Content $M = 6.36$, SD = 2.64). In the five-count conditions, however, more weight was assigned to the cellphone evidence $F(1, 217) = 6.14$, $P = 0.01$, $\eta^2 = 0.03$ (1 Count = 5.69, SD = 2.77; 5 Count = 6.62, SD = 2.80) and also the criminal history evidence $F(1, 217) = 8.79$, $P < 0.01$, $\eta^2 = 0.04$ (1 Count = 4.84, SD = 3.08; 5 Count = 6.06, SD = 2.96). This greater weight on the cellphone evidence in the five-count condition is logical; that the defendant happened to be near five robberies is more indicative of guilt than being near only one.

Those participants finding the defendant guilty and seeking to impose a prison sentence were asked how long that sentence should be (in months). This did not vary significantly across conditions, but there was a non-significant trend in favor of longer sentences in the metadata condition. $F(1,73) = 3.27$, $P = 0.08$, $\eta^2 = 0.04$ (Meta $M = 32.58$, SD = 13.76; Content $M = 26.28$, SD = 13.31). There was no effect of the number of counts and no interaction (both $Fs < 1$).

In addition to the previously described conditions, there was a further variation in which the defendant was again accused of five counts of robbery, but the content information was more explicit. Here, the defendant texted his girlfriend that he had "just knocked over the store on [street of the crime]." The office presenting these texts then explains that "knocked over" can be used as slang for "robbed." Confidence of guilt in this condition ($M = 67.54$, SD = 28.30) was not significantly different from that in the five-count metadata condition ($M = 65.75$, SD = 27.18) or the five-count less-explicit content condition ($M = 64.06$, SD = 28.62). $F(1, 146) = 0.18$, $\eta^2 = 0.00$. There was also no difference in sentence length among those who sought to assign a sentence. $F(2,65) = 2.37$, $P = 0.10$, $\eta^2 = 0.07$ (Meta $M = 33.92$, SD = 12.09; Content-less explicit $M = 26.96$, SD = 13.22; Content-more explicit $M = 26.90$, SD = 13.07).

This study shows that there is not an inherent bias for or against metadata evidence. When content and metadata evidence convey what is logically the same information, respondents will respond appropriately rather than being unduly swayed or turned off by the novelty of receiving metadata in place of content. Also, the lack of a significant difference between the explicit content condition and the less-explicit content condition should not be over-interpreted. As can be seen in Fig. 1, the difference between these conditions is in the expected direction. It was simply too small to be significant. It may be that those respondents who were reluctant to convict in that condition required something more than circumstantial evidence to get beyond a reasonable doubt.

Experiment series 2: sometimes content information gives you something that metadata cannot

In the first study, the content data conveyed the same information as the metadata. In this study, we sought to investigate the ways in which content data could convey increased nuance, particularly for crimes in which state of mind is intensely relevant. We therefore constructed a homicide vignette that gave participants the option of convicting for either first- or second-degree murder. The jury instructions
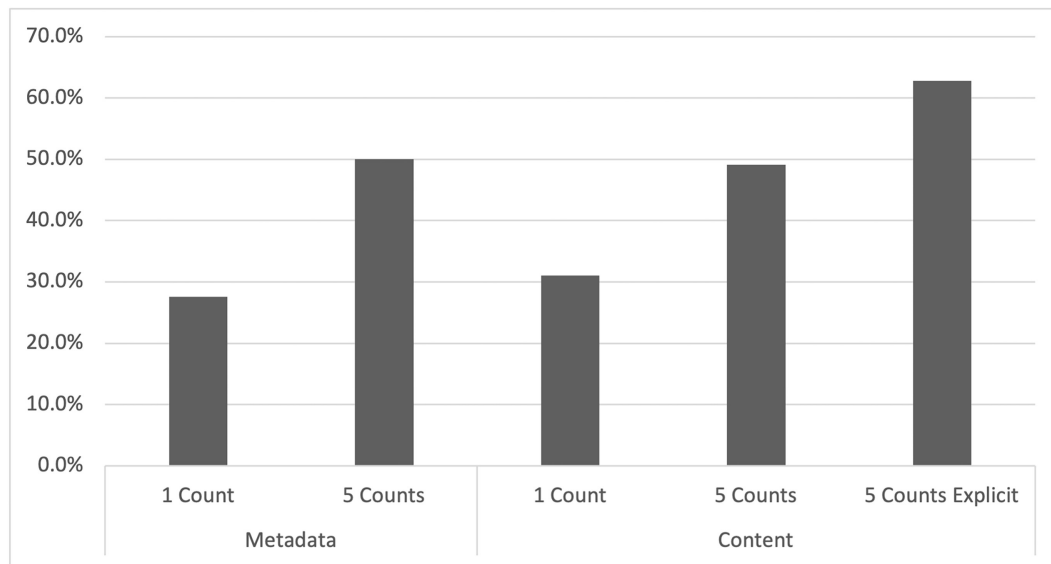
Figure 1: Percentage finding guilt across conditions

provided in this study explained that first-degree murder was murder with premeditation: "This means that the decision to commit murder must be formed before the killing, with enough time to allow for reflection by the defendant." Second-degree murder could occur "unexpectedly or in the heat of the moment." It required only that the defendant "intentionally killed the victim."

The victim of this crime was a woman who had recently ended a relationship with a man. She was described as having been killed in the entryway of her house, her screen door having been forced open. The defendant was her ex-boyfriend. In the metadata condition, an officer testified that he had sent the victim a message from near the victim's house at 9:00 pm, the victim had replied at 9:01, and that the defendant had replied to that at 9:05 pm. The time of the victim's death was believed to be between 9:00 pm and 11:00 pm that night. The full text of this metadata condition is in the Online Appendix.

In the friendly content condition, the messages were said to be as follows:

9:00 pm, Defendant: "I'm here to pick up my stuff. Leave it on the porch."
9:01 pm, Victim: "Sure, one sec."
9:05 pm, Defendant: "Thanks, have a good night."

In the aggressive content condition, they were:

9:00 pm, Defendant: "I'm here. Bring out my stuff before I get mad."
9:01 pm, Victim: "This isn't a good time. Come back tomorrow."
9:05 pm, Defendant: "Don't make me come in there."

No location information was provided in either content condition. These two conditions were intended to provide the range of possible content for the messages. They were either friendly and civil or hostile and aggressive. The other evidence in the case was scant. The defendant's fingerprints were found at the scene, but the officer admits that they could have been weeks old.

The sample was again recruited from CloudResearch. A total of 397 participants passed attention checks. They were 177 men and 220 women. They had a median age of 41 and 42.8% had college degrees, making them somewhat more educated than national norms.

Because of the nature of the crime and the two possible guilty verdicts, participants were separately asked about the probability that the defendant killed the victim and the probability that he did so with premeditation. The probability that he killed the victim varied significantly across conditions, $F(2, 387) = 9.28$, $P < 0.001$, $\eta^2 = 0.05$. Post-hoc tests revealed that participants in the aggressive content condition thought the defendant was more likely to be guilty ($M = 66.26$, SD $= 25.97$) than participants in the metadata condition ($M = 57.53$, SD $= 28.70$, $P = 0.013$) and the friendly content condition ($M = 51.25$, SD $= 29.69$, $P < 0.001$). Participants also thought that the defendant was more likely to have killed the victim in the metadata condition than in the friendly content condition, though this effect was only marginally significant ($P = 0.07$).

Despite this effect on probability of guilt, there was no corresponding difference on probability of premeditated murder, $F(2, 387) = 1.07$, $P = 0.344$, $\eta^2 = 0.01$. This interesting pattern was also reflected in the assigned verdicts. As can be seen in Fig. 2, approximately equal proportions convicted the defendant of first-degree murder across conditions: Metadata $= 21.8\%$, Friendly $= 22.4\%$, Aggressive 19.2%. But the pattern for conviction for second-degree murder was quite different: Metadata $= 24.1\%$, Friendly $= 12.7\%$, Aggressive 36.2%. In short, participants in the friendly condition are less likely to think the defendant killed the victim in a fit of passion but approximately equally likely to think it was part of a cold-blooded plan.

Those participants finding the defendant guilty and seeking to impose a prison sentence were asked how long that sentence should be (in years). This did not vary significantly across conditions. $F(2, 166) = 0.03$, $P = 0.97$, $\eta^2 = 0.00$. On average, people assigned 11.79 years (SD $= 3.43$). Participants did not assign evidence weights in this study.

Experiment series 3: metadata establish patterns
Study 2 was designed to show the power of content data to add color to a set of bare facts. Study 3, in contrast, was designed to show the superior ability of metadata to present pattern evidence. The crime here was a hotel room murder. The victim was described as a married individual whose spouse was out of town. Police were said to suspect that the victim had gone to the hotel to have sex—
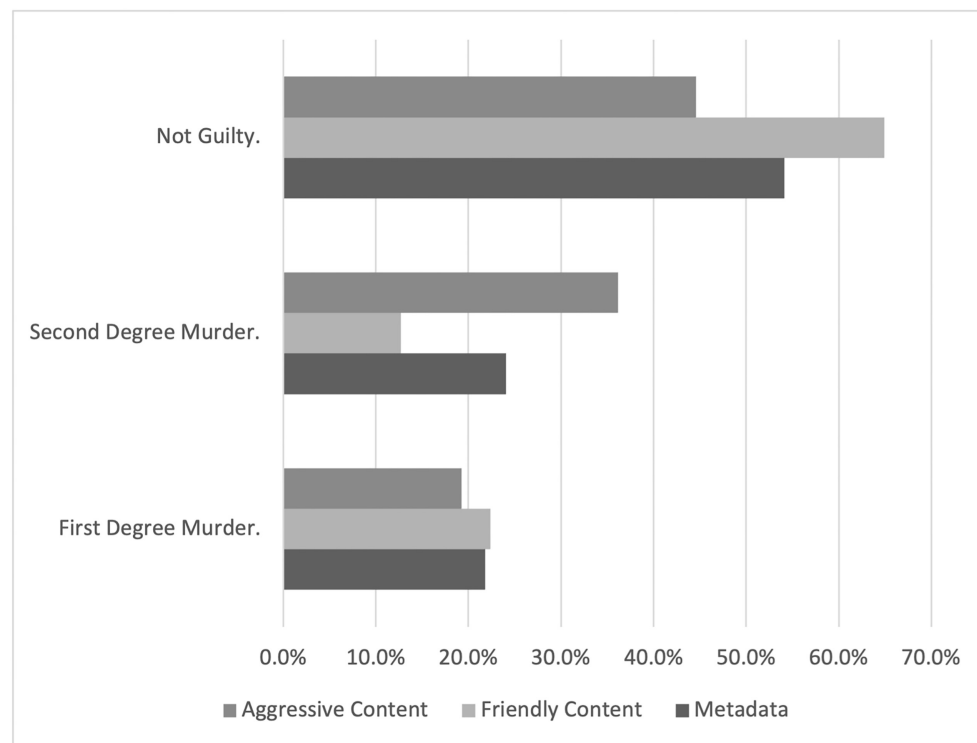
Figure 2: Percentage of people reaching each verdict by condition

condoms and sexual lubricant were present. The defendant was an acquaintance of the victim and their spouse, whom they had met at church.

In all conditions, it was said that the murderer had been observed entering the hotel room by a maid and that the defendant fit the vague description the maid provided. The key evidence was either metadata or witness testimony suggesting that the victim and the defendant had been having an affair. In the strong witness condition, a neighbor testified to having seen the defendant at the victim's apartment complex on several occasions in the evening, including on a particular date when the spouse had been out of town. The witness was said to have remembered the date because their favorite sports team had played a good game that night. The spouse testified to having never been told about any one-on-one visits between the victim and the defendant. The weak witness case was similar, though the witness admitted to not being sure about the date of any particular visit. This was intended to create some ambiguity because the victim's spouse testified to having had the defendant over for dinner on two occasions.

The metadata case incorporated information from a parking-lot gate. People were said to be able to access the apartment complex's parking lot by using an app to open the gate. The defendant, it was explained, had been given access to this app by the victim and their spouse. The access log showed that the defendant had repeatedly visited the apartment on nights when the spouse had been out of town. The full text of this metadata condition is in the Online Appendix.

Since Study 2 had a female victim and a male perpetrator, we were concerned that there might be gender-based patterns in our results. In Study 3, we therefore manipulated whether the victim was male or female, keeping the defendant and spouse as opposite sex. The witness was always male.

The sample was again recruited from CloudResearch. A total of 832 participants passed attention checks. They were 374 men, 455 women, and 3 reporting other genders. They had a median age of 48, and 37.2% had college degrees.

Since the main evidence in this study was directed at establishing the affair between the defendant and the victim—with the implication that the affair partner was also the killer—we asked about both the probability that they were having an affair as well as the probability that the defendant killed the victim. Both of these probability measures differed significantly across evidence conditions. Affair $F(2, 820) = 11.13$, $P < 0.001$, $\eta^2 = 0.03$. Murder $F(2, 820) = 7.25$, $P < 0.001$, $\eta^2 = 0.02$. For the affair, post hoc tests revealed that participants thought the defendant was more likely to be having an affair with the victim in the metadata condition ($M = 73.19$ SD = 23.29, $P < 0.001$) than in the strong ($M = 64.56$, SD = 26.93) or weak ($M = 63.54$, SD = 25.93) witness condition, which did not differ. Similarly, participants thought the defendant was more likely to have killed the victim in the metadata condition ($M = 59.64$, SD = 29.62, $P < 0.001$) than in the strong ($M = 52.77$, SD = 32.03) or weak witness ($M = 50.23$, SD = 31.48) conditions, which again did not differ. As can be seen in Fig. 3, this translated to 48.0% guilty in the metadata condition compared to 39.9% in the strong witness condition and 38.6% in the weak witness condition.

The gender of the participant (coded as male or not for simplicity) had no significant effects or interactions. The gender of the victim had only two significant effects. First, the male victim was judged to be more likely to have been having the affair, $F(2, 820) = 11.57$, $P < 0.001$, $\eta^2 = 0.01$, and this effect was qualified by an interaction between victim gender and condition, $F(2, 820) = 3.56$, $P < 0.05$, $\eta^2 = 0.01$. Simple effects analyses revealed that the male victim was judged to be much more likely to be having an affair in the strong witness condition (12.14 points more likely) than in the weak witness (2.99 points more likely) or metadata condition (2.17 points).

A separate ANOVA examining the effect of condition on sentence length found that, conditional on wanting to assign a prison sentence

Figure 3: Percentage finding guilt across conditions

to the defendant, there was no effect. $F(2, 316) = 1.55$, $P = 0.22$, $\eta^2 = 0.01$. The average sentence was 12.2 years.

As in Study 1, participants here were asked to report how much weight they assigned different pieces of evidence (the connection between victim and defendant through church, the witness or metadata evidence, and the physical similarities between the defendant and perpetrator). The only significant effect was on the witness or metadata evidence. This was weighed more heavily in the metadata condition ($M = 6.25$, SD $= 2.85$) than in the strong witness condition ($M = 5.40$, SD $= 2.97$, $P < 0.001$), and more heavily in the strong witness condition than in the weak witness condition ($M = 4.91$, SD $= 2.80$, $P < 0.05$).[1]

Discussion and policy implications

Our studies show that the relative power of content and metadata information is highly contextual. Content information and metadata can be equally useful when they convey information that is logically equivalent; there is no bias in favor of either (Study 1). Yet content information conveys additional meaning in circumstances where the defendant's state of mind is critical (Study 2). And metadata can more convincingly establish a pattern of behavior (Study 3). Additionally, respondents acting as hypothetical jurors do not appear to recommend longer sentences on average when presented with content information instead of metadata, suggesting that metadata information is not as somehow less emotionally impactful.

These findings have significant implications for shaping the policy conversation surrounding encryption policy. Most notably, our findings suggest that metadata may be stronger evidence for some categories of criminal acts that are frequently foregrounded during conversations about encryption policy—including prosecutions related to criminal organizations, which may rely on patterns of behavior to demonstrate contact between organization members. We conclude this part by discussing the limitations of our study, and the steps we took to minimize them.

---

1 This was rated on a 0–100 scale but is reported on 0–10 to maintain consistency with Study 1.

Content information and metadata are and are not equivalent evidence of guilt

The key characteristics of and results from each study are presented in Table 1. Notably, the three studies differ substantially in the criminal activity being considered, and the role that the metadata/content information played in establishing the facts of the case.

When both content and metadata evidence convey the same type of information—the defendant's location at a particular moment in time—respondents reacted similarly regardless of which they were presented with. One could have expected either the familiarity and richness of message content or the scientific clarity of metadata records to be more persuasive. Instead, they were equivalent in the first study. And both types of evidence showed the same magnitude increase in persuasiveness when the evidence placed the defendant at the scene of five crimes rather than one.

When the content information assigns meaning to what is otherwise an ambiguous pattern, particularly when that meaning bears on a defendant's state of mind, respondents react very differently. In the second study, respondents presented with content information showing that the defendant was at the victim's home and angry around the time the victim was killed were more likely to convict than respondents presented with metadata merely showing the defendant was near the victim's home around the time she was killed. Respondents presented with content information showing that the defendant was at the victim's home and (at least superficially) friendly around the time the victim was killed were less likely to convict than respondents presented with the metadata evidence.

The difference in conviction rates appears to be driven by different assumptions about the defendant's state of mind. In all conditions, respondents were approximately equally likely to convict for first-degree murder. In other words, an equal proportion of respondents in each experimental condition believed that the defendant had shown up to the victim's house intending to kill her. These respondents presumably viewed the "friendly" messages as a cruel deception, written to lull the victim into a false sense of security or trick future investigators. Yet those receiving the two different kinds of content information had sharply different reactions to the question of whether the defendant had killed the victim absent premeditation, or in the heat of passion. Those receiving the friendly messages thought

Table 1: Summary: of study design and results

	Study 1	Study 2	Study 3
Crime	Robbery	First-degree murder (second-degree murder)	Second-degree murder
Role of evidence in trial	Show at location	Show at location	Show preexisting sexual relationship
Metadata case	One or five instances of cell phone being near crime scene at time of crime	Messages sent from near victim's house near time of crime	App showing defendant's car was in parking lot on days when victim's spouse was out of town
Weak content/witness case	Text message reading "just been at the store"	Text message reading "thanks, have a good night"	Witness who saw defendant at apartment complex on several unclear dates
Strong content/witness case	Text message reading "just knocked over the store"	Text message reading "don't make me come in there"	Witness who saw defendant at apartment complex on several dates, including one when V's spouse out of town
Findings—conviction rates	Content information and metadata equivalent under all circumstances; more of each is more powerful	More likely to convict in strong content than metadata; in metadata than weak content	More likely to convict in metadata than witness case; strong witness and weak witness case very similar
Findings—sentence length	Equivalent sentence length for content information and metadata	Equivalent sentence length for content information and metadata	Equivalent sentence length for content information and metadata

this was relatively unlikely. Those receiving the aggressive messages thought this was much more plausible; an angry confrontation was entirely consistent with the tone of what they read. The metadata case sensibly fell between these extremes. Comments from participants in the metadata case showed that people projected a variety of meanings onto the message timestamps, with many saying that they had no idea what the messages might have said.

These results show both the strength and limitations of content information. Content information can shed substantial light on what is otherwise a dark place. Mere message timestamps convey only so much information, and very different pictures of the defendant's intentions are conjured by the differing messages. But even apparently clear messages must be interpreted. When respondents acting as hypothetical jurors are convinced that the crime occurred in a certain way, they may be unpersuaded even when the content information would appear to contradict their understanding of the case.

The second study also demonstrates that the content information lost to encryption and deletion can hurt defendants as well as prosecutors, assuming it is unavailable to both. Even though respondents were more likely to convict when presented with content information showing an aggressive mental state ("don't make me come in there"), they were also marginally less likely to convict when presented with content information showing a friendly mental state ("thanks, have a good night"). A defendant's self-serving testimony about the contents of lost messages would be both risky and subject to doubt, whereas a chat log held by a third party could be introduced without them even taking the stand.

Furthermore, where the evidence in question is used to show a pattern of underlying behavior rather than place the defendant at the scene of the crime, respondents acting as hypothetical jurors appear to be more convinced by metadata rather than by witness testimony. In the third study, respondents were more likely to convict when presented with metadata showing that the defendant had made repeated trips to the victim's apartment at times when the victim was alone than by witness testimony purporting to establish the same facts. One could easily imagine the challenging cross-examination of a witness whose testimony hinged on whether a neighbor was visited by a friend on this or that particular but unremarkable date. Timestamps were better able to establish the critical pattern of visits.

We believe the arc of our results can be at least partially explained by the different types of ambiguity created by metadata and content information. Metadata—especially communications metadata—can very effectively establish the bare facts of an individual's behavior. Person A went here X number of times. Person B sent messages from precisely these locations. But it creates ambiguity about the meaning of that behavior; why were they there, and what were those messages? Content information or witness testimony can remove that ambiguity; those messages were angry, or they were not. But content must also be filtered through individuals—creating new ambiguity about the meaning of communication or the accuracy of remembrances. It will be a rare witness who can recall the arrival times of a car as well as a garage door opener. Which type of evidence is more persuasive depends on which type of fact most needs to be established. Metadata will more easily place the defendant there and then. Content information will more easily show a state of mind, whether exculpatory or inculpatory.

Exposure to content information or metadata may also impact other aspects of juror decision-making. As previously discussed, jurors exposed to content information may be more likely to recommend longer sentences, as content information can provide evidence about the defendant's state of mind at the time of the crime—and, consequently, demonstrate their culpability. However, conditional on respondent wishing to impose a jail sentence, we did not observe a relationship between exposure to content information instead of metadata on the length of the recommended sentence. In the first study, we did not observe a statistically significant difference between recommended sentence lengths when the contents of the defendant's communications related the criminality of their actions ("knocked over a store") as opposed to just the location of the scene of crime ("at the store"), and in fact observed a non-statistically significant trend toward longer sentences when the conviction was predicated on metadata evidence. In the second study, respondents did not appear to recommend different sentences when content evidence showed that the defendant was angry at the time of the crime ("don't make me come in there").

Policy implications

On the whole, our findings present a complicated picture of whether and how metadata may mitigate the impact of encryption on criminal trials, at least in the context of juror interpretation of the available evidence. The rise of metadata appears to make it easier to convince hypothetical jurors to convict where the case in question rests on establishing patterns of behavior while having a more nuanced effect on the rate at which hypothetical jurors vote to convict in cases where establishing the defendant's state of mind is a key component of the prosecution's underlying case. Consequently, we should expect that a shift from presenting content information at trial to presenting metadata at trial—as might be brought about by the simultaneous "going dark" of communications content information due to encryption and the widespread availability of metadata due to the "golden age of surveillance—to have a differential impact on jury deliberations depending on the type of case in question, the type of fact most likely to be in dispute, and the type of agency tasked with gathering the evidence. Notably, however, these differential impacts are not currently reflected in the types of crimes commonly discussed as part of the encryption policy debate.

For crimes where establishing the defendant's state of mind is a critical and difficult-to-accomplish component of the case, content information is a powerful tool that is difficult to replicate through metadata. These crimes vary widely and are likely investigated and prosecuted by a large range of law-enforcement agencies. For example, criminal liability for insider trading under federal law is available only for willful violations of the Securities and Exchange Act [52], which "generally requires the government to show that the defendant acted with the knowledge that the conduct was specifically unlawful" [53]. Content information is uniquely powerful evidence of willfulness in these cases, as it can demonstrate what the defendant understood about their actions at the time they undertook them. Content information may be similarly powerful evidence when prosecuting murder cases, where the difference between obtaining a conviction for first-degree murder or second-degree murder will depend on whether the prosecutor is able to demonstrate deliberation and premeditation on the part of the defendant. Content information may include statements that are relevant to premeditation that would not be obtainable through metadata.

Where the outcome of a trial rests on showing particular patterns of behavior related to repeated communications or visits, metadata appear to provide more convincing evidence than witness testimony. Under these circumstances, metadata has the advantage of appearing both more objective and more reliable, since it was obtained through automatically generated records of the defendant's behavior rather than an individual's observations and recollections. Crimes, where patterns of behavior are particularly powerful evidence of guilt, are likely to become easier to prosecute due to the rise of metadata. For example, prosecutors in cases related to involvement in criminal organizations can now rely on extensive location records to show that the defendant was repeatedly at the same place and time as known members of the organization, suggesting contact.

Rather than seeing current technological trends as causing the information available for criminal investigations and trials to either "go dark" or enter a "golden age," our findings suggest that—from the perspective of jury decision-making—these technological trends will have differential impacts across crimes where prosecutors are likely to rely on different types of evidence. The current debate around policy responses to encryption often does not reflect these differential impacts, however. In a speech on encryption policy, e.g. former Attorney General Barr specifically cited concerns over "transnational drug cartels increasingly mov[ing] their communications onto commercially available encrypted platforms" in order to hide their activity [13]. From the perspective of developing evidence that will convince a jury to convict, our research suggests that this type of crime may be one where metadata may be an effective alternative.

Furthermore, as different types of law-enforcement agencies have different tools available to them to overcome the loss of content information due to encryption, we expect that the "going dark" phenomenon will have the biggest impact on jury deliberations in jurisdictions where law enforcement is least able to respond effectively to encryption. While a variety of "workarounds" to obtain encrypted information have been identified, they are resource-intensive and therefore will be limited only to those agencies that have sufficient resources to bring to bear [54]. Most notably, recent research on the use of surveillance tools by local law-enforcement agencies suggests that many of these agencies are familiar with the use of metadata such as cell phone location information, but that use of highly specialized tools like cell site simulators may be limited to only a few large agencies [55]. These differences in the ability of law enforcement to respond to changes in surveillance technology have broad potential policy and federalism implications, as they may lead to increased reliance on federal prosecutions for certain types of crimes [56]. The policy conversation should focus on those agencies and crimes most likely to be impacted by the loss of content information due to encryption. It currently does not.

Taken as a whole, these findings suggest several concrete steps forward for the encryption policy debate. First, this debate should be centered on cases where the loss of content information is likely to impact criminal investigations and trials. Second, efforts to mitigate the impact of encryption by providing law enforcement with assistance in using other forms of surveillance technology should specifically consider the usability of that technology in smaller jurisdictions with fewer resources. Finally, decision-makers creating policies or tools to facilitate law-enforcement surveillance in response to the investigatory difficulties posed by encryption should consider limiting the use of these policies or tools to types of crimes where law enforcement is least able to use metadata in place of content information. That the lessening availability of content information makes certain crimes harder to prosecute does not mean that a general lessening of privacy, making all crimes easier to prosecute, is the correct tradeoff.

Limitations

There are several significant limitations to this study. First, this study focuses only on the perspective of juror decision-making. However, a criminal inquiry must proceed through many steps before the jury is asked to render a verdict: an investigation (which may include obtaining warrants), the bringing of criminal charges, preliminary hearings (and potentially plea bargaining), and a trial. Each of these stages involves different actors in the criminal justice system, with different incentives and levels of expertise, who may respond to metadata and content information in different ways. While our study speaks directly to how jury decision-making may be impacted by exposure to metadata in place of content information, these results should not be generalized to other stages of the criminal investigation and trial process.

Second, our results are derived from online surveys conducted with individuals who are potentially eligible for jury duty rather than from directly observing the behavior of actual jurors. To the extent that those who are potentially eligible for jury duty acting in a hypothetical context may behave differently from those who are actually selected for and serving on a jury, our study may not accurately

capture real-world behavior. We took several steps to mitigate this limitation. We attempted to roughly demographically match our respondents to the composition of those serving on juries on both age and gender [57], in order to ensure some equivalence between those who participated in our study and actual jurors. However, as it is likely that prosecutors and defense attorneys select jurors on the basis of both observable characteristics (such as age and gender) and unobservable characteristics (such as ability to follow directions and likely sympathy with their arguments), our demographic matching is unlikely to result in a group of respondents that precisely mirrors those serving on juries. Additionally, individuals currently serving on a jury may behave differently from individuals responding to a survey due to contextual differences: the seriousness of being in court, exhortations by the judge, and knowledge that the verdict rendered will substantially impact the defendant's life. We attempted to mitigate these differences by basing our jury instructions on those used in actual criminal trials, in order to mirror the context of a real-world jury as much as possible.

Third, our experiment was not designed to consider the cumulative effect of metadata and content information on juror decision-making. While this decision allowed us to more sharply highlight the potential tradeoffs that occur when different categories of evidence are presented at trial, it also presented somewhat of an artificial dichotomy, as content and metadata are frequently presented together at trial. Future studies into behavioral responses to encrypted devices should consider circumstances where both metadata and content information are available. This is particularly important for studies investigating law-enforcement responses to encrypted devices during criminal investigations, as some studies suggest that metadata may be a complement for content information during criminal investigations under some circumstances [58].

Finally, our experiences while conducting this experiment demonstrate the inherent challenges of investigating whether exposure to the same information conveyed through different types of evidence impacts juror decision-making. During the first study, e.g. we found it challenging to construct text messages that conveyed the same information as location metadata while still being reasonably realistic. This required extensive discussions and decisions about how people communicate in real life: Would anyone really text their partner just to say they were at the store? Would they do it five times? In the real world, location data might put the defendant at the scene of each of five crimes, whereas content information would either only confirm a single location (worse than location data) or be an explicit confession (much better).

Our difficulty in developing scenarios that would allow us to test the equivalence of metadata and content evidence without changing the information established through this evidence showcases both the limitations and value of these experiments. We struggled to present hypotheticals where metadata and content information were used to convey the same information precisely because it is likely that content information would have conveyed information that could not be obtained through metadata, and that metadata may be more consistently available than content information. Consequently, we would expect the extent to which metadata may act as a substitute for content information to depend heavily on the context of the case. Studies 2 and 3 were developed to further explore how varying context might impact the relative strengths of metadata and content information, and ultimately demonstrated the relative strengths and weaknesses of these types of evidence. This research establishes the need for a nuanced understanding of the comparative costs and benefits of presenting content information and metadata at trial when making encryption policy decisions.

Conclusion

In this paper, we set out to inform the conversation regarding whether increased availability of metadata may mitigate impediments to criminal investigations caused by consumer encryption, which in turn is part of a large and fierce debate about policy responses to "going dark." To do this, we provide the first empirical evidence on how exposure to content information or metadata impacts the behavior of any actor in the criminal justice system, focusing on the jurors who are ultimately asked to render a verdict in these cases. Our results suggest that the impact of being presented with metadata instead of content information varies substantially based on the context of the case. While metadata and content information appear to be equally convincing to respondents acting as hypothetical jurors under circumstances where they convey equivalent information, metadata appear to have an advantage when the prosecution is attempting to demonstrate a pattern of behavior and content information appears to have an advantage when the prosecution is attempting to demonstrate something about the defendant's state of mind. Notably, these findings suggest that a shift from content information to metadata brought about by the rise of consumer encryption will have differential effects on different types of crimes—at least when considering the impact of this phenomenon on juror decision-making.

Our results emphasize the need for more evidence about "going dark" to both characterize the scope of this phenomenon and better understand how different actors in the criminal justice system respond to it. Though some of the current encryption policy debate turns on normative questions that cannot be resolved empirically [31], many arguments can be substantially informed by empirical data. Given the substantial potential societal costs and benefits of the encryption policy proposals that have been under consideration in recent years, it is imperative that additional evidence be brought to this debate.

Supplementary data

Supplementary data is available at *Cybersecurity Journal* online.

Acknowledgments

Author contributions

Anne E. Boustead (Conceptualization, Data curation, Fo.

rmal analysis, Investigation, Methodology, Project administration, Validation, Writing—original draft, Writing—review and editing) and Matthew B. Kugler (Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Validation, Writing—original draft, Writing—review and editing)

References

1. Hennessey S. The elephant in the room: addressing child exploitation and going dark. Aegis Paper Series No. 1701, Hoover Institution, 2017.
2. Wainscott A. A "golden key" to Pandora's Box: the security risks of government-mandated backdoors to encrypted communications. *N Ky L Rev* 2017;**44**:57.

3. Kaye D. Report of the Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression (A/73/348). United Nations. 2018. Retrieved from: https://freedex. org/wp-content/blogs (10 February 2023, date last accessed).

4. Balkovich E, Prosnitz D, Boustead A. *et al. Electronic Surveillance of Mobile Devices*. Santa Monica, CA: Rand Corporation, 2015.

5. Solove DJ. A brief history of Information Privacy Law. *Proskauer on Privacy*. PROSKAUER ON PRIVACY, PLI, 2016. GWU Law School Public Law Research Paper No. 215, 2006. https://ssrn.com/abstract=914271

6. Whitten A, Tygar JD. Why Johnny can't encrypt: a usability evaluation of PGP 5.0. *USENIX Security Symposium: Berkeley, CA, USA*, pp. 169–84, 1999.

7. Thompson AW, Park C. *Privacy's Best Friend: The Importance of Encryption in Protecting Consumer Privacy*. Washington, DC: New America, 2020.

8. Farivar C. *Apple Expands Data Encryption Under iOS 8, Making Handover to Cops Moot*. 2014. https://arstechnica.com/gadgets/2014/09/apple-expands-data-encryption-under-ios-8-making-handover-to-cops-moot/ (20 July 2021, date last accessed).

9. Miller J. *Google and Apple to Introduce Default Encryption*. 2014. https://www.bbc.com/news/technology-29276955 (20 July 2021, date last accessed).

10. Rashid FY. *Encryption, Privacy in the Internet Trends Report*. 2019. https://duo.com/decipher/encryption-privacy-in-the-internet-trends-report (20 July 2021, date last accessed).

11. Eckart JP. The Department of Justice versus Apple Inc.—the great encryption debate between privacy and national security. *Cath UJL Tech* 2018;**27**:1.

12. Comey J. *Going Dark: Are Technology, Privacy, and Public Safety on a Collision Course?*, Washington, DC: Brookings Office Of Communications, 2014.

13. Barr WP. Remarks at the International Conference on Cyber Security, 2019. https://www.justice.gov/opa/speech/attorney-general-william-p-barr-delivers-keynote-address-international-conference-cyber

14. Comey JB, Yates SQ. Going dark: encryption, technology, and the balance between public safety and privacy. Joint Statement of Deputy Attorney General Department of Justice and Federal Bureau of Investigation Before the Committee on the Judiciary United States Senate Judiciary Committee, Washington, DC, 2015.

15. Administrative Office of the US Courts. *Wiretap report*, 2019. https://www.uscourts.gov/statistics-reports/wiretap-report-2019

16. Manhattan District Attorney's Office. *Smartphone Encryption and Public Safety*. https://www.manhattanda.org/our-work/smartphone-encryption-and-public-safety/. (August 2021, date last accessed).

17. Lewis JA, Zheng DE, Carter WA. *The Effect of Encryption on Lawful Access to Communications and Data*. Lanham, MD: Rowman & Littlefield, 2017.

18. Swire P, Ahmad K. Encryption and globalization. *Colum Sci Tech L Rev* 2011;**13**:416.

19. Gasser U, Gertner N, Goldsmith JL. *et al. Don't Panic: Making Progress on the" Going Dark" Debate*, Cambridge, MA: The Berkman Klein Center for Internet & Society at Harvard University, 2016.

20. Greengard S. *The Internet of Things*. Cambridge, MA: MIT Press, 2021.

21. Ferguson AG. The Internet of Things and the Fourth Amendment of effects. *Calif L Rev* 2016;**104**:805.

22. Mikusz M, Houben S, Davies N. *et al. Raising Awareness of IoT Sensor Deployments*. London: IET, 2018.

23. Kugler MB, Hurley M. Protecting energy privacy across the public/private divide. *Fla L Rev* 2020;**72**:451.

24. Pell SK. You can't always get what you want: how will law enforcement get what it needs in a post-CALEA, Cybsecurity-Centric Encryption Era. *NCJL Tech* 2015;**17**:599.

25. Ramirez G. What carpenter tells us about when a Fourth Amendment search of metadata begins. *Case W Res L Rev* 2019;**70**:187.

26. Tokson MJ. The content/envelope distinction in Internet law. *Wm Mary L Rev* 2008;**50**:2105.

27. Fuchs C. Web 2.0, prosumption, and surveillance. *Surveill Soc* 2011;**8/3**:288–309.

28. Soghoian C. Caught in the Cloud: privacy, encryption, and government back doors in the Web 2.0 Era. *J Telecomm High Tech L* 2010;**8**:359.

29. Mac R, Haskins C, Sacks B. *et al.* Your local police department might have used this facial recognition tool to surveil you. Find out here. *BuzzFeed News*. 2021.

30. Gartenberg C. Why Apple's new privacy feature is such a big deal. *The Verge*, 2021.

31. Rozenshtein AZ. Surveillance intermediaries. *Stan L Rev* 2018;**70**:99.

32. Zasu Y. Sanctions by social norms and the law: substitutes or complements?. *J Leg Stud* 2007;**36/2**:379–96.

33. Bowles S, Polania-Reyes S. Economic incentives and social preferences: substitutes or complements?. *J Econ Lit* 2012;**50/2**:368–425.

34. Langpap C, Shimshack JP. Private citizen suits and public enforcement: substitutes or complements?. *J Environ Econ Manage* 2010;**59/3**:235–49.

35. Rozenshtein AZ. Wicked crypto. *UC Irvine L Rev* 2018;**9**:1181.

36. National Academies of Sciences, Engineering, and Medicine. *Decrypting the Encryption Debate: A Framework for Decision Makers*. Washington, DC: National Academies Press, 2018.

37. Manhattan District Attorney's Office. *Report of the Manhattan District Attorney's Office on Smartphone Encryption and Safety*. New York, NY: Manhattan District Attorney's Office, 2015.

38. Boustead AE. Small towns, big companies: how surveillance intermediaries affect small and midsize law enforcement agencies. Aegis Series Paper No. 1802, Hoover Institute, 2018.

39. Pennington N, Hastie R. A cognitive theory of juror decision making: the story model. *Cardozo L Rev* 1991;**13**:519.

40. Vorms M, Lagnado D. Coherence and credibility in the story-model of jurors' decision-making: does mental simulation really drive the evaluation of the evidence?. In: *International Conference on Model-Based Reasoning*. 103–19. Berlin: Springer, 2018.

41. Willmott D, Boduszek D, Debowska A. *et al.* Introduction and validation of the juror decision scale (JDS): an empirical investigation of the story model. *J Crim Justice* 2018;**57**:26–34.

42. Klettke B, Graesser AC, Powell MB. Expert testimony in child sexual abuse cases: the effects of evidence, coherence and credentials on juror decision making. *Appl Cog Psychol* 2010;**24/4**: 481–94.

43. Houck MM. CSI: reality. *Sci Am* 2006;**295/1**:84–9.

44. Schweitzer NJ, Saks MJ. The CSI effect: popular fiction about forensic science affects the public's expectations about real forensic science. *Jurimetrics* 2007;357–64.

45. Shelton DE, Kim YS, Barak G. A study of juror expectations and demands concerning scientific evidence: does the CSI effect exist. *Vand J Ent Tech L* 2006;**9**:331.

46. Podlas K. The CSI effect and other forensic fictions. *Loy LA Ent L Rev* 2006;**27**:87.

47. Landau S. Making sense from Snowden: what's significant in the NSA surveillance revelations. *IEEE Secur Priv* 2013;**11/4**:54–63.

48. Penrod S, Hastie R. Models of jury decision making: a critical review. *Psychol Bull* 1979;**86/3**:462.

49. Bornstein BH. The impact of different types of expert scientific testimony on mock jurors' liability verdicts. *Psychol Crime Law* 2004;**10/4**:429–46.

50. Daftary-Kapur T, Dumas R, Penrod SD. Jury decision-making biases and methods to counter them. *Legal Criminol Psychol* 2010;**15/1**:133–54.

51. Chandler J, Rosenzweig C, Moss AJ. *et al.* Online panels in social science research: expanding sampling methods beyond Mechanical Turk. *Behav Res Meth* 2019;**51/5**:2022–38.

52. Securities and Exchange Act, 15 U.S.C. § 78ff. (2020).

53. Moser S, Weitz J. 18 USC Sec. 1348—a workhorse statute for prosecutors. *US Att'ys Bull* 2018;**66**:111.

54. Kerr OS, Schneier B. Encryption workarounds. *Geo LJ* 2017;**106**:989.

55. Oliver M, Kugler M. Surveying surveillance: a national study of police department surveillance technologies (August 25, 2021). 54 Arizona State Law Journal 103 (2022), Northwestern Law & Econ Research Paper No.

21-08, Northwestern Public Law Research Paper No. 21-26, 2021. https://ssrn.com/abstract=3911442

56. Boustead AE. The tools at hand: surveillance innovations and the shifting role of Federal Law Enforcement in Drug Control. *Ohio St J Crim L* 2020;**18**:1.

57. Anwar S., Bayer P, Hjalmarsson R. The role of age in jury selection and trial outcomes. *J Law Econ* 2014;**57/4**:1001–30.

58. Boustead AE. *Police, Process, and Privacy Three Essays on the Third Party Doctrine*. Santa Monica, CA: The Pardee RAND Graduate School, 2016.