The Effect of Violent Video Games on Violent Crime

Robert McDonough*
Gretchen Gamrat†

October 2024

Abstract

We analyze the effect that violent videos games have on violent crime in the United States. Using county-level variation in retail sales of "mature" video games, some of which occur on proprietary platforms, we leverage exogenous variation in exposure to identify corresponding changes in crime outcomes. Especially after high-profile violent crimes, policymakers and the news media frequently argue that increased exposure to violent games will lead to increased violent crime. We find no such evidence. If anything, our analysis suggests that there are short-run decreases in violent crime, specifically violent sexual offenses, following the release of mature video games.

JEL: K42, L82,L86

*UNC Chapel Hill Department of Economics. Email: robmcd@unc.edu
†University of Virginia McIntire School of Commerce. Email: hta9nf@virginia.edu

1 Introduction

Around 40 percent of Americans believe that there is a relationship between violent video games and violent behaviors, and 32 percent of those who play video games believe the same.[1] With the media's support (Markey and Markey, 2015), this belief is understandable. Moreover, given that roughly three out of four Americans play video games—and that the average gamer plays 14 hours per week[2]—any positive causal effect of violent games on violence would be a societal concern. Yet, there is little evidence to support the hypothesis that violent video gaming leads to an increase in violent criminality.

We contribute to this area by examining whether increased exposure to violent video games has an effect on violent crime rates. We measure exposure to violent games at the county level through a dataset that contains video game sales records from a set of retail chains starting in 2007. This dataset allows us to exploit both the timing of violent video game releases and variation in county-level purchase of those releases to estimate the effect that violent video games have on violent crime.

Our identification amounts to an instrumental variables strategy that leverages a technological limitation on the ability to access video games. Namely, video game disks work only when paired with their matching platform. For instance, an Xbox 360 game disk works only on an Xbox 360 console. When *Halo 3*, a violent science-fiction game, was released in 2007 exclusively for the Xbox 360, the only way to play the game was to buy the disk and insert it into an Xbox 360. Since all game disks are coded for a specific platform, a person's ability to access a new game that releases on one or several platforms varies based on the platforms she already owns. By exploiting this hardware-software link, we isolate variation in exposure to newly released games that is exogenous from any unobserved determinants of crime. To implement this identification strategy, we use variation in lagged platform-specific game sales to predict the sale of newly released violent video games. We use these predictions to then estimate the effect that increased exposure to a violent game release has on weekly agency-level violent crime rates obtained from the National Incident-Based Reporting System (NIBRS). We find no evidence that violent game releases lead to increased short-term violent crime rates. Further, the point estimates that we recover are small enough in magnitude to rule out the large effects that have been suggested by some others. We also find evidence suggesting that increased exposure to violent games causes a decrease in the rate of violent sex offenses.

[1] See `https://www.pewresearch.org/internet/2015/12/15/gaming-and-gamers/`.
[2] See `https://www.npd.com/news/press-releases/2020/more-people-are-gaming-in-the-us/`.

The public impression that violent video games *should* lead to violent criminality is informed by a body of social science research on the connection between violent media consumption and aggressive behavior. In laboratory experiments, for example, the effect of violent video game consumption has been studied extensively. In this setting, researchers have found that playing a violent game leads to increased aggression compared to playing a nonviolent video game. The laboratory has allowed researchers to establish, for instance, that violent video games cause an immediate increase in bodily production of stress hormones linked to fight-or-flight responses (Gentile, Bender and Anderson, 2017). To take these findings outside of the laboratory, researchers have also conducted surveys to explore the relationship between violent gameplay and "real-life" behavior (e.g., Möller and Krahé (2009)). These studies find that those who play violent video games more frequently are also more likely to report anti-social and violent behavior, and the reported effect sizes are similar in magnitude to laboratory studies. Collectively, this body of work suggests that the correlation between violent gameplay and aggression is large enough in magnitude to constitute a public health concern, given the number of people playing violent video games. In a meta-analysis of studies of violent gameplay and aggression, Anderson (2003) finds that the correlation between violent video games and aggressive outcomes is similar in strength to the correlation between passive secondhand smoke exposure and lung cancer.

Yet, surveys cannot account for the fact that those who play violent video games in real life have selected into the activity, and are likely different from those who do not select into violent gameplay. For example, Ward (2010) finds that controlling for gender, race, and geography eliminates a positive association between increased video game play and adolescent fighting for all but the most heavy video game users. Beyond demographic variables, other unobserved third factors could easily cause selection into both violent gameplay and violent behavior. The simple fact that video game players *believe* in the cathartic effect of violent gameplay (Olson, Kutner and Warner, 2008) suggests one such confounder. Namely, those with heightened latent aggression turn to video games in an attempt to relieve that aggression, and are also more likely to resort to real violence when no simulated violence is available (e.g., starting a fight in the schoolyard). In the presence of selection like this, correlational studies will not identify the causal effect of playing violent video games on violent crime. Further, the causal effects of violent gameplay on aggression that have been well-identified in the laboratory lack the external validity to inform us about effects on criminality for a variety of reasons. Normally, laboratory researchers are constrained to studying only the immediate effects of exposure to a violent game, as outcomes are normally measured in the minutes or hours following gameplay. Also, laboratory studies cannot evaluate the impact of violent gameplay in the research subjects' normal lives—the setting in which we are concerned about the effect of violent media.[3]

Policymakers' interest in violent video games is predicated on reducing high-impact violent actions.[4] Thus, while research on the correlation between video games and aggressive behavior is a useful starting point, policymakers need evidence on the presence or absence of a causal link between violent video games and violent crimes. We help to provide this evidence by finding and leveraging exogenous variation in exposure to violent video games, following a strand of economic literature that had sought to provide causal evidence on the behavioral effects of popular media. Dahl and DellaVigna (2009), one of the first studies in this area, estimates the short-term effects of violent movies on violent crime. By exploiting variation in the violence of movies shown in a given movie theater across different days, Dahl and DellaVigna (2009) finds that violent movies cause a contemporaneous decline in violent crime rates. In the first causal analysis of the effect of video games on crime in the United States, Cunningham, Engelstätter and Ward (2016) uses game quality and recency of game release as instruments that generate exogenous time series variation in the sale of video games in the United States. This research finds that violent video game exposure caused by new game releases actually leads to small declines in national violent crime rates, but the authors caution that their IV estimates should be viewed more as robustness checks. Indeed, one conclusion the paper draws is that an analysis using cross-sectional variation in game sales—the sort of variation we use—would be a fruitful area for further research.

The identification strategy that we use is also reminiscent of that used in Kearney and Levine (2015) and more recently in Lindo, Swenson and Waddell (2022), both of which examine television media. Kearney and Levine (2015) examines the effect that *16 and Pregnant* had on teen birth rates, while Lindo, Swenson and Waddell (2022) examines the impact that *The Ultimate Fighter*—a reality TV show about mixed martial arts fighting—had on violent crime rates. Both of these studies rely on the notion that television viewers exhibit habit persistence with regard to the TV stations they watch. When a new show premiers, there is variation in exposure to that show driven by how many

[3] Researchers also debate how to interpret the results of laboratory experiments *vis-à-vis* "aggression." For instance, some argue that outcome measurements used in laboratory studies (e.g., whether a participant chooses to conclude an open ended story prompt with a violent or non-violent resolution) have an unclear link to an aggressive state of mind. Ferguson (2007b) provides a good overview of this and other common criticisms.

[4] See, for instance, the opening statements made during the original 1993 U.S. Senate hearing on the potential regulation of violent video games: https://www.govinfo.gov/content/pkg/CHRG-109shrg28337/html/CHRG-109shrg28337.htm.

people were already in the habit of watching a certain channel.[5]

While some of the same habit persistence likely does exist for the platforms that gamers use to access their games, our identification strategy does not rest on this behavioral assumption. Rather, we are exploiting the technical fact that gamers can access a new video game only if they have the correct hardware. Thus, when a new video game releases for one or more platforms, an individual's actual exposure to that game is determined partially by whether they already own the appropriate hardware. The video game sales that we observe are platform-specific (i.e., we can distinguish between an Xbox disk and a PlayStation disk), meaning that we can use lagged platform-specific game sales as a proxy for the stock of gaming platforms in a county. When new violent games then release for a certain set of platforms, counties are differentially exposed to those games based on their prior stock of hardware. Also, since we can measure the national sale of game releases in our data, we can leverage differences in the popularity of specific games—and differences in a game's popularity between platforms—to increase the precision of our estimates.

Our results do not show that increased exposure to new violent video games after their release is associated with any increase in violent crime outcomes, as measured by weekly crime rates reported by police agencies. In most cases, we estimate treatment effects that are precisely centered on zero. With regard to violent sexual offenses, we actually find that increased exposure to violent video games leads to statistically significant *decreases* in these crimes. More striking than the statistical significance of any individual result is the fact that our point estimates are consistently small in magnitude. Our estimates of the effect of additional exposure to violent game releases on assaults, for instance, are generally on the order of 0.1 percent of a standard deviation in assault rates.[6]

Using an event-study specification, we then ask whether the effect of violent video game releases varies across weeks. Estimating the week-by-week effect of game exposure in this way helps us to assess whether the null results from our pooled model are obfuscating any important time-varying effects. For instance, other researchers have pointed out that *incapacitation*, the mechanical fact that video game play displaces time which could be used for alternative activities, could exist alongside other mechanisms. Effectively, this implies that any point estimate for the effect of video game exposure on crime could reflect an agglomeration of both incapacitation effects (negative) and psychological effects

[5] Lindo, Swenson and Waddell (2022) actually goes further than this by leveraging the fact TV viewers turn to specific channels during specific time slots. We summarize both of these studies in more detail in Section 2.

[6] The standard deviations in crime rates that we refer to are calculated cross-sectionally across agencies for a given time period. We describe our standard deviation calculation in detail in Section 3.

(uncertain). We find no evidence for any differences in treatment effects across weeks using our event study models, suggesting that the precise null results from our pooled estimator reflect the absence of any causal relationship between violent video games and violent crime.

Since we exploit changes in the ability to play new violent video games over the short run among populations that have already expressed an interest in video gaming, our analysis is most well-equipped to isolate effects along the intensive margin of violent game playing– we find no effect on violent crime of an increase in playing violent games by communities that are already playing violent games to some extent. Our results are largely silent with regard to the effect of changes in violent gaming along the extensive margin—the effect of communities (and individuals) playing violent games over long periods of time, compared to the counterfactual in which those communities are not playing violent games. This implicates our ability to detect phsycological mechanisms in particular, as these could manifest slowly and over longer time horizons than we can consider.

2 Background

2.1 Violent media, behavioral outcomes, and crime

In 1993, the Senate held the first congressional hearings on violent video games. Of particular concern to Senators was the marketing of violent games to children, the increasing realism of violence available through games, and the lack of cohesive regulation of violent game sales. The consequence of the 1993 Senate hearings on video game violence was the voluntary adoption and enforcement of a content rating system administered by the Entertainment Software Rating Board (ESRB). Akin to the Motion Picture Association of America (MPAA) responsible for the American movie rating system, the ESRB rates games on an escalating scale based on content. Games that have pervasive and realistic violence are rated *M: for mature*, which the ESRB defines as being "generally suitable for ages 17 and up. May contain intense violence, blood, and gore."[7] Video game retailers aere supposed to sell M-rated games only to those above the age of 17, and participating retailers are supposed to verify customer age with photo ID.

The 1993 Senate's concern over violent video games was predicated on the longstanding theory

[7] The two other common rating categories are *E: for everyone*, for games that are entirely suitable for all ages; and *T: for teen*, for games with some violence but minimal blood and no extreme violence. See www.esrb.com for a complete history and description of the content rating system.

that engagement with violent media leads to an increase in violent behavior, including violent crime. Theories on the relationship between media and aggression can be traced to antiquity, but the general aggression model (GAM) (Anderson and Bushman, 2002a) is the predominant contemporary theory used to explain how something like violent gameplay would lead to aggressive behaviors such as violent crime. The GAM can be traced back to Bandura's social learning theory (Bandura, 1977), but also incorporates refinements such as script theory (Heusmann, 1988), cultivation theory (Gerbner et al., 1994), and others. In this framework, violent video games provide users with repeated opportunities to participate in simulated violent behavior that is rewarding and satisfying, leading gamers to develop internal scripts in which aggression is a viable course of action (Anderson and Dill, 2000), and also teaching players to expect hostile and aggressive actions from others (Anderson and Bushman, 2002b).

Empirically, the GAM has largely been tested through laboratory work. Laboratory experiments are generally conducted by randomizing participants to play either a violent or non-violent video game for a period of time, and shortly thereafter measuring some outcome variable linked to aggression.[8] Proponents of the GAM point to these lab studies as cohesively showing that video games lead to relatively high levels of increased aggression. In a meta-analysis, Anderson (2003) finds that across multiple aggression measurements, playing violent video games leads to a 0.26α increase in aggression, which the author notes is a larger effect size than that found for secondhand cigarette smoke and lung cancer. However, Ferguson (2007b) and other critics have raised concerns about whether the outcomes used in these studies are well suited to measure aggression, as well as how to interpret the magnitude of these treatment effects. Overall, though, the claim that violent video games have some short-term impact on aggression in the laboratory is well supported.

In order to take the GAM's predictions about video games outside the laboratory, researchers have largely relied on correlation analyses in cross-sectional and panel surveys of adolescents (e.g., Möller and Krahé (2009), Anderson et al. (2008) and Gentile et al. (2004)). While these studies find that higher exposure to violent games is associated with increased aggression and aggressive behavior, critics have pointed out that these research designs fail to account for potential problems such as omitted variable bias, or selection into violent gameplay. Indeed, many such surveys do not report how the correlation between violent gameplay and aggression changes after controlling for covariates

[8] Commonly used outcomes meant to measure aggression include how a subject chooses to complete an open-ended scenario prompt (Anderson and Bushman, 2002b), the choice to act aggressively or punitively against opponents in simple games (Anderson and Benjamin, 2004), and even physiological markers linked to flight-or-fight responses such as cortisol levels and cardiovascular arousal (Gentile et al., 2017).

known to be associated with aggression and gameplay, such as gender. Using the CDC's Youth Risk Behavior Survey, Ward (2010) finds that controlling for demographic characteristics of youth in the sample leads the observed correlation between video game play and increased fighting propensity to shrink and become insignificant for groups who play up to four hours of video games daily. Even for those adolescents who report playing over 5 hours of video games *daily*—the group for whom the observed correlation between gameplay and fighting is strongest—Ward (2010) finds that increased fighting propensity relative to those who do not play games is reduced from 13.4 percent down to only 6 percent with the addition of these covariates. Ward argues that even if this remaining positive association is interpreted as causal, the magnitude of the effect is likely not large enough to merit policy intervention. Ferguson (2007a) also notes the descriptive fact that youth-involved crime rates have fallen steadily across the United States during the same period of time in which violent video games have become massively popular.

In contrast to the GAM, other theories of aggression and behavior predict that violent video games would reduce acts of aggression and violent crime. Some researchers have proposed that violent video games reduce violent behaviors by providing gamers with a simulated environment in which to sate aggressive urges, in what is commonly referred to as a theory of catharsis. Konečni and Doob (1972) provides a treatment of modern theories of catharsis that is useful when considering video games by distinguishing a mechanism of catharsis through displaced aggression that is applicable to video games. By providing an outlet through which aggression can be expressed virtually, violent video games relieve the gamer's urge to engage in violence against some real-life cause of aggressive feeling, such as a bully. Notably, many violent video game players believe in the cathartic effect of these games. In a survey of eighth-grade boys who play violent video games, Olson et al. (2008) finds that one common reason why adolescents choose to play violent games is to relieve feelings of aggression, or even to displace a specific urge to act violently in real life. Empirical research, however, finds limited evidence for catharsis effects. Kersten and Greitemeyer (2021) argues that video game players who report catharsis are conflating a general improvement in mood following violent gameplay with actual reductions in aggression.

Finally, violent video games could also serve to reduce crime through a mechanism of incapacitation, which draws on the basic notion of time use popularized by Becker (1965). Entertainment media is generally time-consuming, and violent video games are no exception. When an individual chooses to spend time consuming violent media, they are substituting time away from other activities. In the most straightforward story, some individuals drawn to violent behavior may be directly substituting time away from violent real-life activity in order to spend more time with the virtual violence of a game. More generally, video games would create an incapacitation effect whenever a game player chooses to spend time gaming that would otherwise be spent in an activity with a higher chance of leading to aggression. Indeed, even non-violent games could create incapacitation effects, so long as the time spent playing the game is not being substituted for time spent on an equally non-violent activity. Empirical work focusing on the incapacitating effect of video games is limited, though Ward (2018) finds that popular video games cause gamers to reduce their school attendance during the period when they are completing the game.

While these theories provide potential mechanisms through which violent video games could affect violent crime, well-identified causal analysis of this question is rare. Cunningham, Engelstätter and Ward (2016) is one of the first studies to focus on this causal relationship, using variation in violent game exposure created by new game releases to identify the impact of violent video games on national crime rates, finding evidence for either a null effect or actually a negative relationship between exposure to violent media and violent crime outcomes. Suziedelyte (2021), on the other hand, finds cross-sectional variation in violent video game exposure at the individual level through the timing of interview dates for families in the Panel Study of Income Dynamics (PSID). PSID interview dates are random throughout a given year, meaning that some families are interviewed soon after the release of major violent video games, while others are not. Suziedelyte (2021) finds that adolescents more recently exposed to violent video games before their PSID interview were less likely to engage in violent behavior, though their parents were more likely to report destructive behavior not against people (e.g., damaging school property).

2.2 Modern video game hardware and software

By the 21st century, the video game industry had become a major component in U.S. entertainment spending. In 2008, 53 percent of American adults reported playing video games. By 2020, that share had risen to 75 percent. As the video game market grew, so too did the number of new video games released each year. The online entertainment database IMDb.com contains a listing of 824 new games released during 2010, a 25 percent increase from the number of games released in 2000. Violent video games make up a sizeable share of these new releases. Between 2007 and 2011, roughly one quarter of the best-selling games each year were given the mature content rating, as measured by VGchartz.com, a video game research firm that publishes yearly lists of the top-100 video game releases. Based on VGchartz.com's list, a top-selling mature game was released every 3 weeks on average. As the industry grew and video game releases became more numerous, video game publishers also became increasingly sophisticated with regard to the marketing, product differentiation, and strategic timing of video game releases. Engelstätter and Ward (2018) finds that major video game publishers strategically choose the release date for their new games in order to avoid competition with other games of a similar genre or ESRB rating.

Modern video games need to be understood as a combination of software and hardware. Video games themselves are software developed by video game publishers. A video game can be played only when that software is run on a compatible video game platform, where a platform is defined as any piece of computing hardware and operating system that can play a video game. Given the need for game platforms, the video game market is used as a modern example of a two-sided market (Davidovici-Nora and Bourreau (2012), Rysman (2009)), in which platform manufacturers serve as the intermediary between consumers and software producers. In the early 2000s, common video game platforms were the personal computer, the video game console (e.g., Xbox 360, PlayStation 3), handheld gaming devices (e.g., the Gameboy), and even mobile phones.

As with any other software, video games must be coded to function on specific platforms, meaning that game publishers must choose in advance a set of platforms on which their games will work. When consumers choose to buy a video game, they must then choose a platform-specific version of that game; a game disk designed for one platform cannot be used to play on another, though gameplay is largely identical across platforms. For example, *Call of Duty 4: Modern Warfare* was one of the top game releases in 2007, and was released on PC and on most major game consoles, but was not originally released on any handheld device or on the Nintendo Wii. A consumer who bought *Call of Duty 4* for the Xbox 360 would be able to use the disk to play on any Xbox 360 console, but not on PS3 or a PC.

Most games are released on multiple different gaming platforms. Releases are synchronized in time across both platform and geography, and video game price is also synchronized across software versions. In fact, there is little variation in price even between games sold by competing publishers. Occasionally, a video game is released as an "exclusive" for one platform or family of platforms. For instance, new games in the *Halo* franchise, a series of violent science-fiction games, are published by Xbox Game Studios and released exclusively for the latest Xbox consoles. When *Halo 3* was released in 2007, anyone with an Xbox 360 could buy and play the game, while anyone without access to an Xbox 360 could not.

Video game consoles—computers designed specifically to run video game software in the home—were second only to personal computers as the platform of choice during the early 2000s. Video game consoles are expensive, long-term purchases that remain functional to play the latest video games for roughly a decade. For the past several decades, three major firms have dominated the video game console market: Sony, Microsoft, and Nintendo. In November 2005, Microsoft launched the Xbox 360, while in November 2006 Sony released the Playstation 3 (PS3) and Nintendo released the Wii. Together these three are referred to as the 7$^{\text{th}}$ generation of consoles, and these consoles remained dominant in the video game hardware market until the next generation of consoles was released starting in 2012. 7$^{\text{th}}$ generation consoles were wildly successful products, and led to technological advances including the rapid proliferation of online console gaming. In 2008, 53 percent of adult gamers and 89 percent of teenage gamers used a console at least some of the time. At release, the Xbox 360 sold for \$400, the PS3 sold for \$500, and the Wii sold for \$250.[9] Console manufacturers generally price their hardware at a significant loss that is then recouped via software licensing fees. As a consequence of this pricing strategy, console manufacturers compete to grow a large and loyal user base that will purchase enough games over the lifetime of the console to ensure profitability (Williams, 2002).

Those who play video games usually possess more than one platform with which to access video game software overall (Davidovici-Nora and Bourreau, 2012); however Derdenger (2014) finds little evidence that gamers purchase more than one console of the same generation. While some gamers are drawn to one console or another for idiosyncratically reasons (e.g., the desire to play a specific exclusive available on only one console), consoles are effective substitutes, since each provides access to hundreds of new video games each year. Moreover, while some exclusives are popular, many non-exclusive games are equally or more popular. And specifically with regard to violent games, many violent exclusive and non-exclusive games were available for both the Xbox 360 and the PlayStation 3.

[9] These were the prices for the baseline or standard console version sold at launch. Both the PS3 and the Xbox 360 were available in several models at release, with different models largely distinguished by hard drive size and the appearance of the console exterior.

3 Data

We measure video game sales at the county level using the Nielsen Retail Scanner Data (scanner data) spanning from 2007 through 2011. The scanner dataset contains weekly pricing and volume data for products sold in over 35,000 participating stores, comprising roughly 90 retail chains and capturing roughly 30 percent of mass merchandise sales and 50 percent of food and drug sales. In participating stores, Nielsen measures total weekly sales volume and price at the product level using the universal product code (UPC) system. The scanner data contains reported sales for over 2 million distinct UPCs, which are grouped into roughly 1,100 product categories. In the "video and computer games" product category, we have retail sales records for 22,328 UPCs. With regard to video games, it is useful to think of UPCs as each representing a distinct instantiation of a product. For example, *Call of Duty 4: Modern Warfare* was one of the most popular games released in 2007. Not only was it initially released on several consoles, but it was also re-released several times due to its popularity. Nielsen's scanner data contains about a dozen distinct UPCs corresponding to a game disk for *Call of Duty 4*. Of those, one UPC corresponds to the original game disk used to play *Call of Duty 4* on the Xbox 360, while another corresponds to the equivalent game disk used for the PlayStation 3.

We use giantbomb.com, an online wiki and database for information about video game software and hardware, to identify the set of video games released from 2007 to 2011, as well as to capture product information such as release dates and content ratings. Using this information, we manually identify the set of UPCs associated with each video game from this time period that was released for either the Xbox 360, the PlayStation 3, or the Nintendo Wii. We successfully identify at least one UPC code for 84.5 percent of the 2,329 games that released on at least one of the three consoles noted above. Out of the roughly 500 games for which we found no sales records, the majority were either international games that saw no major U.S. release, or games that were released only via digital storefronts (e.g., the Xbox Live Arcade Game Pack, a collection of arcade games released directly on Microsoft's digital storefront). While each UPC corresponds to a console-specific version of a game, the product descriptions included in Nielsen's scanner data do not usually indicate this detail. To identify the console associated with each UPC, we rely on upcitemdb.com, an online UPC database. Out of the 22,000 video game UPCS for which Nielsen has sales records, we match over 16,000 of those UPCS to a specific console.

Not all of the games released during our sample period saw widespread sale in the United States.

In order to identify the video game releases that would be popular and accessible across the U.S., we also scrape yearly video game popularity data from VGchartz.com. We use these top-selling game lists to identify mature video game releases that were major enough to be used as treatment events. In Table 1, we provide a complete list of the top-selling mature game releases between 2007 and 2011. Within a given year, there is a sizeable gradient in sales between the games on VGchartz's bestsellers list. In 2007, for instance, VGchartz lists Wii Sports as the best-selling game, with 6 million game units sold in that year in the US. For comparison, VGchartz reports that the 100th best-selling game in 2007 sold only 250 thousand copies. In Table 2, we break out the proportion of top-selling games that received each ESRB content rating, as well as the fraction of those top-selling games for which we see sales records in the Nielsen panel. Of the games on VGchartz.com's lists of top releases, roughly one quarter are rated M. We find sales records for 80 percent of the games on the top-100 lists, and 97 percent of M-rated games.[10] A large majority of the top-selling games for which we do not see game sales are E-rated games released exclusively for handheld video game platforms (e.g., Pokemon Diamond/Pearl Version, released only on the Nintendo DS).

In Figure 1, we show the time series of weekly mature video game sales revenue across our panel, plotted against the release dates for top mature releases. Visualizing the data this way, we can see that sales of mature games at retail stores in our dataset spike following the release of top-selling mature games.

To get a sense of how much of the total U.S. video game market is present in Nielsen's sample of retailers, we calculate the total monthly and yearly video game sales revenue present in the scanner data. Using national sales figures published by the NPD Group as the denominator, we estimate that Nielsen's scanner dataset captures between one and two percent of new physical video game sales. The most important factor driving this low figure is that physical sale of new video games in this time frame was heavily concentrated among a handful of mass-market retailers. Nielsen censors the name and exact location of all physical stores, and the data-sharing agreement with Nielsen prohibits attempts to uncover retailer names, as well as any mention of specific retailers that are present or absent from the dataset. Speaking generally, though, it is logical to conclude that the handful of retailers specializing in video game sales are not present in Nielsen's scanner data, since the presence

[10] VGchartz ranks games separately by console, so for instance *Call of Duty 4* is in the top-100 game list for 2007 twice, once for the PS3 and once for the Xbox 360. This means that the number of unique games in each top-100 list is less than 100. From 2007 to 2011, there are 285 distinct games in VGchartz rankings, of which we find sales records for 226.

of any one of these retailers would lead to a higher observed fraction of total video game sales. This supposition also explains the lack of console sales in the scanner data. The normal profit margin on video game platforms like the Xbox 360 is far lower than that of video game disks. Thus while many stores sell video game disks, fewer general merchandise chains find it profitable to stock video game hardware.

Using scanner data, we produce county-by-week sales figures for each video game UPC observed. For each video released for either the Xbox 360, the PS3, or the Wii, we then calculate weekly sales figures by console for the first eight weeks after each video game release. We also aggregate across different video games and calculate the weekly video game units sold for each console in each county. Since video games must be purchased for a specific platform, these console-by-county software sales allow us to proxy for the hardware being used in each county, a decision we discuss in more detail when we outline our empirical strategy. Finally, we calculate the total weekly revenue spent on all products in the video and computer games category for each county.

We use the National Incident-Based Reporting System (NIBRS), accessed via the Inter-university Consortium for Political and Social Research at the University of Michigan, as our primary source for crime data. Intended as the successor for the Uniform Crime Reports (UCR), NIBRS is administered by the FBI and provides information on criminal activity at the incident-level. For a criminal incident, NIBRS records information about the date and time of the incident, the offenses committed, the victims, and (where known) the offenders. Crimes are reported into NIBRS by the law enforcement agency with jurisdiction where the incident occurred, meaning that weekly crime patterns can be constructed at the agency or county level. NIBRS offense codes are quite granular, including 46 specific crimes against either person (e.g., homicide), property (e.g., motor vehicle theft), or society (e.g., illegal gambling). However, participation in NIBRS by law enforcement agencies is not universal. As of 2011, only 36 states had law enforcement agencies that reported into NIBRS, either indirectly through a state reporting program or directly to the FBI. Only 32 percent of those agencies that report into the UCR reported into NIBRS in 2011. Despite these limitations, NIBRS provides the only data source with which weekly crime patterns can be analyzed at a granular geographic level, and has been increasingly used by empirical researchers. We aggregate NIBRS offense-level data to create agency-by-week counts for the offense categories included in NIBRS. NIBRS also includes an agency-by-year population estimate for the area served by each agency, allowing us to express our crime outcomes as rates *per capita*. This agency population variable also lets us reliably exclude law enforcement agencies with an unclear geographic jurisdiction, such as state patrols and university police departments.

We link the county-by-week sales figures for video game releases from Nielsen scanner data with the agency-by-week crime counts from NIBRS to construct a panel dataset with observations at the week-by-agency level. The law enforcement agencies reporting into NIBRS are nested within counties, meaning that county-level sales figures from Nielsen can potentially be matched to multiple agencies. However, we follow Lindo et al. (2022) and do not aggregate our dataset to the county level for our primary analysis. Some law enforcement agencies report data into NIBRS only sporadically, or stop reporting into NIBRS in certain years, or cease reporting into NIBRS altogether during our sample period. Thus, conducting our analysis at the agency level ensures that changes we observe in crime outcomes are not driven by changes in the agencies reporting into NIBRS. However, given that the exogenous variation we exploit is at the county level, we cluster our standard errors at the county level for our primary analysis.[11] Given that we observe only one to two percent of yearly video game sales, there are also many small counties for which we observe negligible sales of any product in the video game category. Likewise, since police agencies report data into NIBRS by incident, a lack of reporting by a given agency for a period of time could reflect sporadic reporting or simply an absence of crime, particularly when an agency is in a sparsely populated county. To mitigate the issue of sparse reporting of either video game sales or crime, we restrict our primary analysis to agencies that report a positive population within their jurisdiction that are nested within a county with a population of at least 100,000 overall. Our final dataset includes roughly 1500 police agencies spread across 200 counties.

4 Empirical Strategy

Our identification strategy takes advantage of the fact that playing a video game requires a specific combination of hardware and software. Once consumers have purchased one video game platform, they are more exposed to new video games that are available for that platform. For instance, when a violent video game comes out for the Xbox 360, individuals who own an Xbox 360 are more likely to purchase and play the Xbox 360 version of that game. Due to the requirement that software must match hardware, the consumers who are more exposed to one of these game releases are those who

[11] See Maltz and Targonski (2002) for a discussion on the challenges of using county-level crime data.

already own one of the video game platforms on which the new release can be played.

Since our available data covers the period of time in which the 7th generation consoles were released and became dominant in the U.S. video game market, we choose to focus on the Xbox 360, the PS3, and the Wii[12]. While we focus on these game releases for at least one of these three consoles, many of these games were also released on some other platforms, and most commonly on personal computers (e.g., *Call of Duty 4* was also released for PC and Mac). To account for this in a feasible way, we combine all other platforms into a single "other" category when measuring video game purchases. We leverage variation in take-up of these four platforms across U.S. counties as a source of exogenous variation in subsequent exposure to violent video games, asking whether counties that should be more exposed to new violent games due to pre-existing hardware popularity experience a relative change in crime rates after those violent games release. This identification strategy is similar to that used by Kearney and Levine (2015), in which the authors instrument for exposure to the television program *16 and Pregnant* using MTV viewership across the U.S. from before the show's premiere date, asking whether increased exposure to *16 and Pregnant* led to a change in teen pregnancy outcomes.

For this strategy to identify the causal effect of violent video game releases on crime, however, variation in console popularity can influence crime outcomes only by influencing the extent to which different counties are exposed to new video games. If selection into purchasing the Xbox 360, the PS3, or the Wii is in any way correlated with crime outcomes, then the exogeneity assumption of our instrumental variables strategy would be violated. This is akin to the issues raised in Jaeger et al. (2016) and Jaeger et al. (2020) regarding the identification strategy used to estimate the effects of *16 and Pregnant* in Kearney and Levine (2015)—that counties with higher MTV viewership before the show premiered were different than counties with lower viewership with respect to trends in teen pregnancy. Here, we can likewise imagine that counties with a large stock of Xbox 360 hardware prior to the release of a violent game for the Xbox 360 are systematically different with respect to crime outcomes than counties with fewer Xbox 360s. For this reason, we identify off of variation in historic console-specific sales *controlling for overall video game spending per capita* in each county.[13] The thought experiment we then envision is that two locations have an equal propensity to consume

[12] While we include mature game sales for the Wii in our analysis, we see far fewer sales of such games on the Wii. This is due in large part to the fact that fewer mature games were released for the Wii compared to other consoles, as Nintendo focused on marketing the Wii to families and "non-gamers."

[13] Given the relative rarity of most violent crimes, we express crime outcomes and local video game sales as rates per 10,000 residents in most of our analysis. For simplicity in our main text, we refer to all such measurements as *per capita* rates.

video games, but vary in terms of the hardware medium with which residents consume those games. When a game like *Halo 3* comes out for only the Xbox 360, counties with an equal demand for violent games in general will still vary in their actual demand for *Halo 3*, solely because of variation in the popularity of the Xbox 360. This is similar to the identification strategy used in Lindo et al. (2022) to estimate the effect on violent crime of the TV show *The Ultimate Fighter*, a popular mixed martial arts fighting show that premiered on Spike TV. Since Spike TV marketed itself as "the first network for men" and tried to target a young male audience, one could easily imagine that general viewership of Spike TV correlates with violent crime rates. To overcome this possibility, this research actually controls for overall Spike TV viewership and exploits only variation in Spike TV viewership during the specific timeslot when *The Ultimate Fighter* was broadcast. Rather than needing to assume that there is no selection into Spike TV overall, this refinement means that Lindo et al. (2022) needs only assume that there is no selection into specific timeslots of Spike TV programming between counties that view Spike TV at similar rates.

By using only the variation in exposure to violent games coming from platform variation, we likewise loosen the exogeneity assumption under which our estimated treatment effects are well identified. Specifically, we must assume that the choice to use one console versus another is uncorrelated with crime outcomes in counties that have a similar level of video game spending overall. Our identifying assumption would be violated, for instance, if counties in which the Xbox 360 became more popular had different crime patterns than counties that played video games to a similar extent, but on the PS3. Since we have data on crime rates both before and after each video game releases, we can relax this assumption further by comparing changes in crime rates before and after violent games release. The validity of our analysis then relies on the assumption that crime trends are parallel between counties with varying console stock.

Most of the video games that we study were released on several platforms, meaning that there are likewise several channels through which a county could be more or less exposed to a given violent video game. Consequently, we can also recast our identification strategy as employing a shift-share or Bartik instrument (Bartik, 1991). The canonical example of a Bartik instrument, as in Blanchard and Katz (1992), uses local employment shares by industry alongside national employment growth rates by industry. Interacting these two variables creates a prediction for local employment growth rates by industry that can instrument for actual employment growth rates. In our setting, we use local console popularity alongside national platform-specific sales of violent video game releases to similar effect. When a video game releases for a given set of platforms, a county's exposure to that game is determined by the share of the county using those platforms to play video games, and the national popularity of that game on that platform (i.e., the shift). Framed as a shift-share instrument, the identification assumption required by our research design is unchanged: the pre-existing hardware popularity in counties must be conditionally exogenous from trends in crime rates (Goldsmith-Pinkham et al., 2020).

While the intuition for our identification strategy relies on the relationship between hardware stock and software sales, as noted in Section 3 we do not observe a usable quantity of hardware sales. Rather, we are limited to observing a consistent fraction of video game software sales. But, game disks function only on one platform, meaning that consumers always purchase a hardware-specific version of a game. Since Nielsen's scanner data lets us distinguish between game disks sold for different platforms, we can use platform-specific software sale in a county prior to the release of a video game as a proxy for stock of that platform at the time of game release. In plain language, we assume that if we see a county largely purchasing games for the Xbox 360, for instance, this is reflecting the fact that gamers in that county are largely playing games on the Xbox 360. In doing so, we must assume the same conditional exogeneity of software sales for a platform as we did for hardware stock of that platform.

This assumption further highlights that we must define the period over which we measure console popularity. Since we use a dataset with multiple game releases in our first stage, one option is to construct a game-specific instrument using lagged sales from one or several months prior to the release of each game. However, this lagged-sales instrument is arguably more vulnerable to violations of the exclusion restriction, especially given our reliance on software sales to infer hardware stock. For example, between 2007 and 2011 three new games were released within the *Halo* franchise. The *Halo* games released for Xbox 360 were popular and violent, and one could imagine that over time a subset of the population also predisposed to violent crime began to select into the Xbox 360 console to take advantage of violent opportunities like *Halo 3*. That is, prior software sales over time could reflect selection into some consoles correlated with violent crime. Such selection would mean that using a lagged prior sales instrument admits the very type of endogeneity that would have biased an estimator using raw game sales.

Alternatively, one could consider using as a prior sales instrument only those sales that took place directly after the release of the new generation of consoles. As argued above, consumers who buy a new console are likely to continue buying games for that console long after the purchase. When the 7th generation of consoles was first released, the hit video games that might induce selection into one console or another were in the future for the most part, and unknown to consumers. Thus, a fixed prior sales instrument is more likely to satisfy the exogeneity requirement. However, as the gap in time between prior video game sales and new violent game releases grows, the predictive power of a first stage using a fixed instrument would naturally wane. This highlights a crucial tension in this identification strategy: by increasing the distance between our prior sales period and our "treatment events" we feel more confident in our exogeneity assumption being satisfied, but must accept a weaker first stage.

Consequently, we run our primary analysis using two versions of our instrument. First, we use a fixed prior sales instrument, using all game sales from the year 2007. The Xbox 360 released in late 2005 and the PS3 and Wii in late 2006, making 2007 the first year in which the entire 7th generation of consoles was available for purchase. Thus 2007 was the period of time during which many consumers were making the decision regarding which new console to purchase. Second, we use three months (12 weeks) of lagged sales, leaving out the one month (4 weeks) of sales directly prior to each game release. For example, since *Halo: Reach* released in September, 2010, the lagged sales instrument would be constructed using county sales data from May, June, and July of 2010. We adopt this "leave-one-out" strategy with our lagged sales instrument specifically because we are using software sales to proxy for hardware sales. Prior to the release of a major video game, individuals who already possess a certain console may change their purchasing behavior (e.g., by saving their money for a month instead of buying a game). Thus, by leaving out the month directly prior to game releases, we assuage concerns that any such anticipatory behavior influences our first-stage predictions. Since we use sales from 2007 as one of our instruments, we run our first stage using release sales for all video games released between 2008 and 2011, including 60 M-rated video games. In Section 5 we show evidence for the validity of our first-stage design, as well as the tradeoff we face in terms of first-stage power.

One low-hanging argument for the violation of our parallel trends assumption comes from price difference between new consoles. While each of the 7th generation consoles were all relatively expensive, their prices at launch were not identical. Such price differences could have led counties to select into different consoles based on socioeconomic factors, which could easily lead to the type of endogeneity we describe above.[14] Consequently, we also control for county socioeconomic factors (i.e., poverty rates, unemployment rates, and per capita personal income) in our analysis.

To implement the first stage of this instrumental variables strategy, we use the county-level sales data available through Nielsen to predict how well newly released violent video games will sell in different counties as a function of those counties' prior platform-specific software sales. Prior software sales and release sales of new games are measured in units (i.e., game disks) *per capita*. In our second stage, we consider the effect of violent video game releases over time periods of varying length, asking how violent game releases impact violent crime in the week following release, in the two weeks following release, and so on. Thus, for each of these time periods we estimate our first stage using sales of the newly released game over the same period. For example, when estimating the effect that a violent game release has on violent crime rates in the two months after that game releases, our first-stage model predicts the total sale of that video game in the two months following its release. We also highlight that two violent video games occasionally release during the same week. To account for this, our dependent variable measuring release sales also aggregates across all of the mature games that release during the same week. That is, we predict total sales in county c over a number of weeks w of the violent video games that release during week t:

$$\begin{aligned} ReleaseSales_{c,w,t} =& \alpha_1 \sum_{g \in G_t} B_{c,g,w,t} + \alpha_2 VideoGameRevenue_{c,t} + \\ & \alpha_3 X_{c,t} + \delta_{w,t} + \mu_{c,t}, \end{aligned} \tag{1}$$

regressing total sales on video game revenue *per capita*, socioeconomic controls, period fixed effects $\delta_{w,t}$, and on our shift-share instrument for a given violent game $B_{c,g,w,t}$. Since multiple games can release for different consoles during the same week, we sum our shift-share instrument across G_t, the set of game releases in week t. For a given game g, our shift-share instrument takes the form

$$B_{c,g,t} = \sum_p PriorSales_{c,p,t} NationalSales_{p,g,w,t,-c}$$

where $NationalSales_{g,p,t,-c}$, our measurement of the national popularity of a game g on a platform p, is the national total sales figure for that game, on that platform, over the same number of weeks. We exclude a county's own game sales when we construct these national sales estimates, hence the $-c$

[14] Indeed, the higher price of the PS3 at launch ($500) compared to the Xbox 360 ($400) has been described as one of the reasons that the Xbox 360 was much more successful at launch than the PS3.

subscript.

in the second stage, we then use the variation in predicted total sales of new violent video games released in week t to identify the effect of violent video game releases on crime. We first fit two-period difference-in-differences models to estimate the impact of violent video game sales on violent crime rates. As noted above, one advantage of this two-period model is that we can vary the window of time around violent video game releases—asking what is the overall effect of an increase in exposure to violent video games in the week following a game release, the two weeks following that release, and so on. The proposed mechanisms through which violent video games impact crime would very likely operate at different points in time after those games release. Any incapacitation effect of violent game releases are liekly to be concentrated in the initial weeks after game release, when most individuals are playing the new game. The increased aggression such as that predicted by the GAM could occur at any point after individuals begin playing the game, but would be more difficult to detect when incapacitation through gameplay was still occurring. An estimated treatment effect close to 0 in the single week following violent game releases could thus reflect a true null relationship between violent gameplay and violent crime, or an offsetting combination of incapacitation and increased aggression—individuals playing the new game are more inclined to violent crime, but have less time to pursue it. Later, we turn to event-study designs in order to estimate week-by-week treatment effects. By first pooling the weeks spanning the release period using a parsimonious difference-in-differences design, however, we can chart any evolution in the aggregate effect of violent game releases on crime while retaining more precision. Our second-stage difference-in-differences equation takes the form

$$\Delta y_{ac,w,t} = \beta_t \Delta Total\widehat{ReleaseSales}_{c,w,t} + \delta_{w,t} + \varepsilon_{ac,w,t}, \tag{2}$$

where $\Delta y_{ac,t}$ is the change in the rate of crime *per capita* in a given agency nested within a county, in the w weeks after one or more mature games release in week t. By expressing the model in differences, agency fixed effects and all of our covariates are differenced out.[15] The $\hat{\beta}_t$ estimate obtained from this model would then be interpreted as the change in crime *per capita* in a police agency's jurisdiction associated with one more violent game sale *per capita* in the corresponding county.

Next, we use an event study to estimate the week-by-week effect of violent game releases on violent

[15] Given that this pooled model has only a pre-period and post-period, a differences model is econometrically equivalent to a two-way fixed effects model. We make the choice to use a differences specification to avoid the computational burden of calculating fixed effects for thousands of police agencies.

crime. In turning to an event-study specification, our goal is to consider whether the release of one or more violent video games has a changing effect on crime rates over time. Given the frequency of mature games releases, however, there are few opportunities from 2007 to 2011 to observe "long" periods over which there are no releases, then the release of a mature game, then a "long" period of time before another mature game is released. Rather, the release of most violent video games will be compounded in the following weeks by the release of another violent game. We can avoid the challenges associated with making inference when treatment is compounding in this way by focusing on the rare mature game releases that are isolated in time from other mature releases. For example, the longest period of weeks in the time series during which only one week contains a mature game release is plus/minus five weeks—this occurs twice, with the release of *Grand Theft Auto IV* (on the Xbox 360 and the PS3) and *Devil May Cry 4* (on the Xbox 360, the PS3, and Windows computers). But there is also a larger set of games that we can use to estimate the release-week effect of violent games more precisely—there are 36 mature games that were released with no other mature releases in the one week plus/minus. Thus, by re-estimating the effect of violent game releases on crime using stratified samples for which there are progressively longer periods of weeks over which only one mature game releases, we can retain more precision in our event study specification. In so doing, we identify off of more game releases when we consider smaller intervals of time (around release dates) than when we consider longer periods of time. Making sample restrictions of this kind, we estimate the following event-study specification:

$$\begin{aligned} y_{ac,t} = &\sum_{j=-J}^{J} \beta_j (WeeksFromRelease = j) \widehat{TotalSales}_{ac,4,t} + \\ &\sum_{j=-J}^{J} \theta_j (WeeksFromRelease = j) VideoGameRevenue_{c,t} + \\ &X_{ac,t} + \delta_t + \lambda_{ac} + \varepsilon_{ac,w,t} \end{aligned} \tag{3}$$

where J is the number of weeks over which there is no release of another violent video game and $j = -1$ (i.e., the week before a game releases) is used as the reference period. In our primary analysis, we use the estimated sale of new games during the month of their release as our measure of county exposure to new violent games. Notice also that the games with longer periods of isolation are subsets of those games with shorter periods of isolation. That is, since no mature games released in the plus/minus five weeks surrounding *Grand Theft Auto IV* and *Devil May Cry 4*, we can use these games when we estimate the release week effects of violent games. With all this in mind, we estimate this event-study specification for $J \in \{1, 2, 3, 4, 5\}$.

5 First-stage Relevance and Evidence for Validity

In Figure 2 we visualize the strength of our first stage in predicting the sale of just-released violent video games. As described in Section 4, we run this first stage separately to predict the release sales of violent games over periods of increasing length, ranging from one week to two months. Here we present the results from the models predicting release month sales, but our first-stage results are similar across all of the release window lengths that we consider. In Figure 2 we plot our instrument on the x-axis and the actual figures for the sale of newly-released violent video games during their release month on the y-axis. Recall that our instrument is the interaction of national game popularity (measured in total units) and county-level platform exposure (measured in units of platform-specific software *per capita*). In Figure 2 Panel A (I), we plot the relationship between our instrument and our release window sales. As we discuss above, we seek to exploit only the variation in exposure to a new violent video game arising from differences in platform popularity between counties that have similar demand for video games overall. In Panel A (II), we plot the same relationship after adjusting for the controls used in our first-stage specification. In both of these panels, there is a strong, visible relationship between our instrumental variable and our first-stage outcome.

We also present the results of our first-stage regression in Table 3. In Panel A, we present the first-stage results obtained using 2007 software sales as our instrumental variable. As shown in Panel A, the first-stage relationship between our shift-share instrument and sales of just-released violent games is statistically significant and positive, even with the addition of both our county-level video game spending control and socioeconomic controls. The corresponding F-statistics are well beyond traditional thresholds used to asses first-stage strength. In Panel B of Table 3, we likewise present the first-stage results using the three-month leave-one-out version of our instrument described in Section 4. Using this alternate version of our instrument produces similar results, though the magnitude of the relationship between our instrument and new violent game sales is larger. This result makes sense, given that this alternate first stage is using sales data that occurs closer in time to the violent game releases we study.

The F-statistics in Table 3 demonstrate that our shift-share instrument is highly predictive of release-month sales for new violent video games. Beyond first-stage relevance, however, we are also interested in evidence for the validity of our research design, which relies on the argument that exposure to violent video games on a specific console should be modulated by the extent to which county residents have previously purchased that console. In order to get a better sense of how this variation in platform popularity contributes to the strength of our instrument, we also run an alternative specification of our first stage, where we replace the estimates of national (platform-specific) game popularity in Equation 1 with platform indicator functions, which turn on if a game is released for a given platform. Our instrument $B_{c,g,t}$ then takes the form

$$B_{c,g,t} = \sum_{p} PriorSales_{c,p,t} * \mathbb{1}_p\{NationalSales_{p,g,t} > 0\}$$

This specification of our first stage allows us to ask whether the story we seek to tell about our instrument is believable: that individuals are more exposed to new video games if they already possess the hardware to play those games.

Figure 2 Panel B and Table 3 Panel B show the results from this alternate first-stage specification. The most striking aspect of Panel B is that, while our first-stage relationship remains strong overall, there are two distinct patterns between 2007 software sales and new violent game sales. Specifically, there is a strong correlation between 2007 sales and new violent sales for one subset of the data (i.e., above the dotted best fit line), while there is almost no such correlation for another subset of the data. We argue that this clearly illustrates the useful role that national game popularity plays when used in our first stage. As noted in Section 3, the release-period sales for the best-selling game in a year are orders of magnitude higher than the corresponding sales for the hundredth best-selling game in a year. But without the estimates of national popularity, the model shown in Panel B is forced to estimate the relationship between prior hardware popularity and new violent game sales as if all of those games are equally important. The consequence of this is that the model underestimates the relationship between hardware popularity and new sale of the most popular games, and overestimates that relationship for the least popular games. Said differently, the choice to interact platform-specific game sales in 2007 with the national popularity of new games allows our first-stage model to weight how important each game release is, yielding a better estimate of the relationship between pre-existing hardware popularity and new software sales.

Panel B also demonstrates that, while including national game popularity improves the strength of our first stage, the platform-specific software sales that we employ are independently useful to predict new game sales. We interpret this as evidence that our instrumental variable is working due to the channel that we hypothesize: counties in which a given platform is already popular are more exposed to new game releases that happen for that console. Moreover, in Panel B (II), we again show that controlling for overall game spending in a county (along with socioeconomic variabes) does not eliminate the predictive power of our instrument.

We have also argued that controlling for overall interest in video games is important to our exclusion restriction. Counties that have more interest in video games generally likely also have more interest in new video game consoles, and a general interest in video games could be correlated with the unobserved determinants of crime. We want to only take advantage of variation in console popularity between counties that have similar interest in video games overall. In Figure 3, we show some evidence for this argument. In Panel A, we show that the software sales *per capita* we observe in our dataset for the Xbox 360 and the Playstation 3 are indeed systematic across the United States. Sales *per capita* for both of these consoles are heavily concentrated in the Pacific Northwest. In Panel B, we then show the residual variation in console-specific game sales that is left after adjusting for video game spending and socioeconomic controls. Panel B demonstrates that controlling for *per capita* spending on video games in 2007 eliminates this systematic variation in console popularity. We argue that this provides suggestive evidence for our exclusion restriction—after controlling for general interest in video games, the variation in exposure to specific consoles that we exploit is not systematic across the United States.

6 Results

Having shown evidence for the relevance of our first stage, and also some suggestive evidence for the validity of our exclusion restriction, we now turn to the results from our difference-in-differences and event-study specifications, which we consider to be our primary results. We focus on the four crime categories included in NIBRS that involve the use of force or violence: homicides, assaults, violent sexual offenses, and robberies.

In Figure 4 we plot the estimated effect of increased violent game sales on violent crime from the two-period difference-in-differences version of our identification strategy. For each crime outcome, we begin by estimating the aggregate effect of increased violent game sales on crime during the two months following a game release week, and tighten the number of weeks included in the post-period (and pre-period) window in one-week intervals. Plotting our results across tightening post-period windows in this way allows us to start thinking about the relative strength of the mechanisms through which violent video games could affect violence. To make our results more interpretable, we present all of our primary results in terms of effect size or impact. Given that our identification takes advantage of variation in game exposure across space, we use the average standard deviation in 2007 crime rates across agencies to calculate effect sizes.[16]

Our difference-in-differences results provide several main takeaways. First, across violent crime categories and post-period windows we find no evidence of any statistically significant positive relationship between increased violent game sales and violent crime. In fact, when we consider a post-period window of between three and six weeks, increased violent video game sales lead to a statistically significant decrease in the rate of violent sexual offenses committed (Panel B). The size of this decrease is roughly 0.5 percent of a standard deviation in violent sex offenses, on average. Though not statistically significant, the estimated effect of violent game sales on homicides (Panel C) and robberies (Panel D) are also negative. The estimated effect of violent game sales on assaults (Panel A) is negative when we consider any post-period between one and five weeks, and positive thereafter, though all of the point estimates for the relationship between mature game exposure and assault are statistically insignificant.

More important than the statistical significance of any individual point estimates, our difference-in-differences estimation allows us to rule out any systematic patterns or empirical regularities running from violent video game release to violent crime. The largest positive treatment effect we estimate is for violent game sales and assaults in the seven weeks after game releases, where we find that one additional violent game sale *per capita* residents increases assaults by .1 percent of a standard deviation over the two months following a game release. Figure 4 also shows that the corresponding 95 percent confidence intervals are tight enough to rule out positive effects larger than 0.6 percent of a standard deviation. Indeed, we rule out positive effect sizes greater than 1 percent across all post-period windows and violent crime categories. In Figure 5, we express the same results as impacts, again showing that violent video game releases are not associated with any large increase in any type of violent crime. Taken together, these results suggest a small incapacitation effect in the weeks immediately following violent game releases.

[16] Specifically, for the difference-in-differences estimator with a post-period of n weeks, we first calculate the standard deviation in crime rates across agencies for each n-week period in 2007. We then average across the standard deviations for each period. In the model with a post-period of just one week, for example, this amounts to calculating the average weekly standard deviation for a given crime type across agencies.

We also estimate the same difference-in-differences model using lagged platform-specific software sales, as opposed to the equivalent sales from 2007. We show the results from this alternate specification in terms of effect size in Figure 6 and in terms of impact in Figure 7. Our results are largely unchanged using this alternate instrument. Again, we estimate treatment effects that are largely negative, and only statistically significant for violent sexual offenses.

Next, we consider the results from our event studies. In moving to an event-study specification, we are no longer estimating the aggregate effect of video games on violent crime over some post-period window, but are instead estimating week-by-week treatment effects. While this gives us more of an ability to tease out any dynamics or patterns in the impact of violent games over time, this ability comes at the cost of precision. For reasons we describe in Section 4, we estimate our event-study specification using only some mature game releases, specifically those that are isolated in time from other mature releases. This means that we lose precision not only because of the increased numbers of parameters that we estimate, but also due to the smaller available sample. In spite of this, the value of our event-study specification is that we can use it to check for any patterns or systematic trends in the week-by-week effect of mature game releases. In our difference-in-differences specification, we can estimate only the aggregate effect of a game release over some post-period window. A point estimate from this specification centered on zero—as in Panel A of Figure 4 showing the effect on assaults—could thus reflect either a true null relationship between violent video games and violent crime, or some offsetting combination of incapacitation and increased aggression. While they do not amount to a formal test of this possibility, the week-by-week effects that we can estimate in an event-study allow us to examine the evolution of crime rates after a game release for patterns that are suggestive of some mechanisms. If the point estimates for the release-week effect of mature games on assaults were negative but then trended towards zero in subsequent weeks, for instance, this would be consistent with a story of incapacitation.

However, the results from our event-study specifications do not reveal any such patterns. In Figure 8 we show the results from this event-study analysis of violent game releases on assault rates. We again present our results in terms of effect sizes. In Panel A, we show the results for the set of games that are isolated in time from other mature releases by at least one week plus/minus. In Panel B, we likewise show the results from estimating our event-study model for the set of games that are isolated for plus/minus two weeks, and so on for Panels C, D, and E. While we see more noise and wider confidence intervals around these point estimates, no visible patterns or trends emerge. Rather, these panels again show point estimates that hew close to 0 and are statistically insignificant.

In Figure 9, Figure 10, and Figure 11, we visualize our event-study results for our other violent crime categories, sex offenses, homicides, and murders, respectively. Across these violent crime categories, we again find no evidence of a positive relationship between violent video games and violent crime. Moreover, while the confidence intervals around our estimates are larger than those from the difference-in-differences estimator, the point estimates themselves remain small for the most part. Further, while some of our point estimates do move away from zero, this is largely happening as we enforce a longer period of isolation from other mature game releases in order to estimate treatment effects further out from release weeks. In Figure 10 Panel E, for example, the estimated effect on homicide in week 1 after the release week of a violent game is almost 0.1α. While this corresponds to a large effect, note that the confidence interval for this point estimate contains the confidence interval for the equivalent week 1 treatment effect in Panel B—recall also that the set of games informing the point estimates in Panel E is a subset of that in Panel B. This suggests that the large point estimate for week 1 in Panel E is simply a consequence of the heavy restrictions we have made on the dataset. When we use all of the available data to estimate the week 1 treatment effect, we see that the more precise estimate turns out to be close to zero.

Overall, the results from our event-study analysis do not reveal any systemic patterns in the impact of violent game releases in the five weeks following the release of such games. Combined with the results from our pooled difference-in-differences estimator, we find no evidence to suggest that violent video games lead to increases in any type of violent crime in the weeks after those games release.

7 Conclusion

We find no evidence that violent video game releases lead to increases in any type of violent crime. Our estimation strategy takes advantage of the fact that the video game hardware that gamers own influences their ability to access new violent video games. Effectively, we compare places that have similar interest in video games, but that vary in their ability to access new video games due to variation in the popularity of different game platforms in those places. With access to cross-sectional variation in video game sales, we are able to leverage this intuition to identify the short-run effects of violent video game releases on violent crime.

For all types of violent crime, we fail to recover any significant positive impact of mature game releases on violent crime. Indeed, our only statistically significant results suggest that mild *decreases* in violent sexual offenses occur after the release of violent video games. We also use an event-study specification in an attempt to tease out whether these null results could be the result of offsetting positive and negative causal channels between mature game releases and violent crime, but again fail to find any evidence for a relationship between violent crime and exposure to mature game releases.

We believe that our results are most useful in conjuction with other studies that have examined the impact of violent media exposure on violent crime using well-identified empirical strategies, including studies on the effect of violent movies (Dahl and DellaVigna, 2009), violent television (Lindo, Swenson and Waddell, 2022), and violent video games (Cunningham, Engelstätter and Ward, 2016). Taken together, these studies and our own provide a solid body of evidence that exposure to violent media does not lead to an increase in violent behavior, and in fact could lead to short-lived decreases in crime. More generally, this growing literature should invite skepticism of the common claim that video games contribute to violent crimes. Short-run effects of violent gameplay on aggression have been documented in a plethora of laboratory studies, but the evidence we present here calls into question how those results translate to outcomes outside of the laboratory. As video games continue to grow in popularity, calls for the regulation of violent games are unlikely to subside. We suggest that policymakers exercise caution before devoting resources to interventions rooted in a fear that violent video game play will lead to real-life violence.

2.8 Acknowledgements

Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

References

Anderson, Craig A., "An update on the effects of playing violent video games," *Journal of Adolescence*, 2003, *27*, 113–122.

— **, Akira Sakamoto, Douglas A Gentile, Nobuko Ihori, Akiko Shibuya, Shintaro Yukawa, Mayumi Naito, and Kumiko Kobayashi**, "Longitudinal effects of violent video games on aggression in Japan and the United States," *Pediatrics*, 2008, *122* (5), 1067–1072.

— **and Arlin J Benjamin**, "Violent Video Games: Specific Effects of Violent Content on Aggressive Thoughts and Behavior," *Advances in Experimental Social Psychology*, 2004, *36*, 199–249.

— **and Brad J. Bushman**, "Human aggression," *Annual Review of Psychology*, 2002, *53* (1), 27–51.

— **and** — , "Violent Video Games and Hostile Expectations: A Test of the General Aggression Model," *Personality and Social Psychology Bulletin*, 2002, *28* (12), 1679–1686.

— **and Karen E. Dill**, "Video Games and Aggressive Thoughts, Feelings, and Behavior in the Laboratory and in Life," *Journal of Personality and Social Psychology*, 2000, *78* (4), 772–790.

Bandura, Albert, *Social Learning Theory*, Prentice-Hall, 1977.

Bartik, Timothy J., *Who Benefits from State and Local Economic Development Policies?*, W.E. Upjohn Institute, 1991.

Becker, Gary, "A theory of the allocation of time," *Economic Journal*, 1965, *75* (299), 493–517.

Blanchard, Olivier Jean and Lawrence F. Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, *1.*

Cunningham, Scott, Benjamin Engelstätter, and Michael R. Ward, "Violent Video Games and Violent Crime," *Southern Economic Journal*, 2016, *82* (4), 1247–1265.

Dahl, Gordon and Stefano DellaVigna, "Does Movie Violence Increase Violent Crime?," *The Quarterly Journal of Economics*, 2009, *124* (2), 677–734.

Davidovici-Nora, Myriam and Marc Bourreau, "Les marchés à deux versants dans l'industrie des jeux vidéo (Two-sided markets in the video game industry)," *Réseaux*, 2012, *173-174*, 97–135.

Derdenger, Timothy, "Technological tying and the intensity of price competition: An empirical analysis of the video game industry," *Quantitative Marketing and Economics*, 2014, *12*, 127–165.

Engelstätter, Benjamin and Michael R. Ward, "Strategic timing of entry: evidence from video games," *Journal of Cultural Economics*, 2018, *42*, 1–22.

Ferguson, Christopher J., "Blazing Angels or Resident Evil? Can Violent Video Games Be a Force for Good?," *Review of General Psychology*, 2007, *14* (2), 68 – 81.

— , "The Good, The Bad and the Ugly: A Meta-analytic Review of Positive and Negative Effects of Violent Video Games," *Psychiatric Quarterly*, 2007, *78*, 309–316.

Gentile, Douglas A., Patrick K. Bender, and Craig A. Anderson, "Violent video game effects on salivary cortisol, arousal, and aggressive thoughts in children," *Computers in Human Behavior*, 2017, *70*, 39–43.

Gentile, Douglas A, Paul J Lynch, Jennifer Ruh Linder, and David A Walsh, "The effects of violent video game habits on adolescent hostility, aggressive behaviors, and school performance," *Journal of Adolescence*, 2004, *27* (1), 5–22.

Gerbner, George, Larry Gross, Michael Morgan, and Nancy Signorielli, *Growing up with television: The cultivation perspective*, Lawrence Erlbaum Associates, Inc, 1994.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Heusmann, L. Roweel, "An informationprocessing model for the development of aggression," *Aggressive Behavior*, 1988, *142*, 13–24.

Jaeger, David A., Theodore J. Joyce, and Robert Kaestner, "Does Reality TV Induce Real Effects? On the Questionable Association Between 16 and Pregnant and Teenage Childbearing," *IZA Working Paper 10317*, 2016.

_ , _ , and _ , "A Cautionary Tale of Evaluating Identifying Assumptions: Did Reality TV Really Cause a Decline in Teenage Childbearing?," *Journal of Business & Economic Statistics*, 2020, *38.*

Kearney, Melissa S. and Phillip B Levine, "Media Influences on Social Outcomes: The Impact of MTV's 16 and Pregnant on Teen Childbearing," *American Economic Review*, 2015, *105.*

Kersten, Riccarda and Tobias Greitemeyer, "Why do habitual violent video game players believe in the cathartic effects of violent video games? A misinterpretation of mood improvement as a reduction in aggressive feelings," *Aggressive Behavior*, 2021, *48* (2), 219–231.

Konečni, Vladimir and Anthony N. Doob, "Catharsis through displacement of Aggression," *Journal of Personality and Social Psychology*, 1972, *23* (3), 379–387.

Lindo, Jason M., Isaac D. Swenson, and Glen R. Waddell, "Effects of violent media content: Evidence from the rise of the UFC," *Journal of Health Economics*, 2022, *83.*

Maltz, Michael D and Joseph Targonski, "A Note on the Use of County-Level UCR Data," *Journal of Quantitative Criminology*, 2002, *18*, 297–318.

Markey, Patrick M. and Charlotte N. Markey, "Violent Video Games and Real-World Violence: Rhetoric Versus Data," *Psychology of Popular Media Culture*, 2015, *4* (4), 277–295.

Möller, Ingrid and Barbara Krahé, "Exposure to violent video games and aggression in German adolescents: a longitudinal analysis," *Aggressive Behavior*, 2009, *35* (1), 75–89.

Olson, Cheryl K, Lawrence A Kutner, and Dorothy E Warner, "The Role of Violent Video Game Conent in Adolescent Development," *Journal ofAdolescent Research*, 2008, *23* (1), 55–75.

Rysman, Mark, "The Economics of Two-Sided Markets," *Journal of Economic Perspectives*, 2009, *23*, 125–143.

Suziedelyte, Agne, "Is it only a game? Video games and violence," *Journal of Economic Behavior and Organization*, 2021, *188*, 105–125.

Ward, Michael R, "Video Games and Adolescent Fighting," *The Journal of Law & Economics*, 2010, *53* (3), 611–628.

Ward, Michael R., "Cutting class to play video games," *Information Economics and Policy*, 2018, *42*, 11–19.

Williams, Dmitri, "Structure and Competition in the U.S. Home Video Game Industry," *The International Journal on Media Management*, 2002, *4*, 41–54.

Figure 1: Weekly revenue from Mature game sales and Mature game release dates (2007-2011)

We plot weekly revenue (in thousands of dollars) from the sale of "mature" rated video games against the release dates for mature games in VGchartz.com's yearly top-selling list. Each red line represents one such mature game release. We observe spikes in mature game sales following the release dates for best-selling mature games

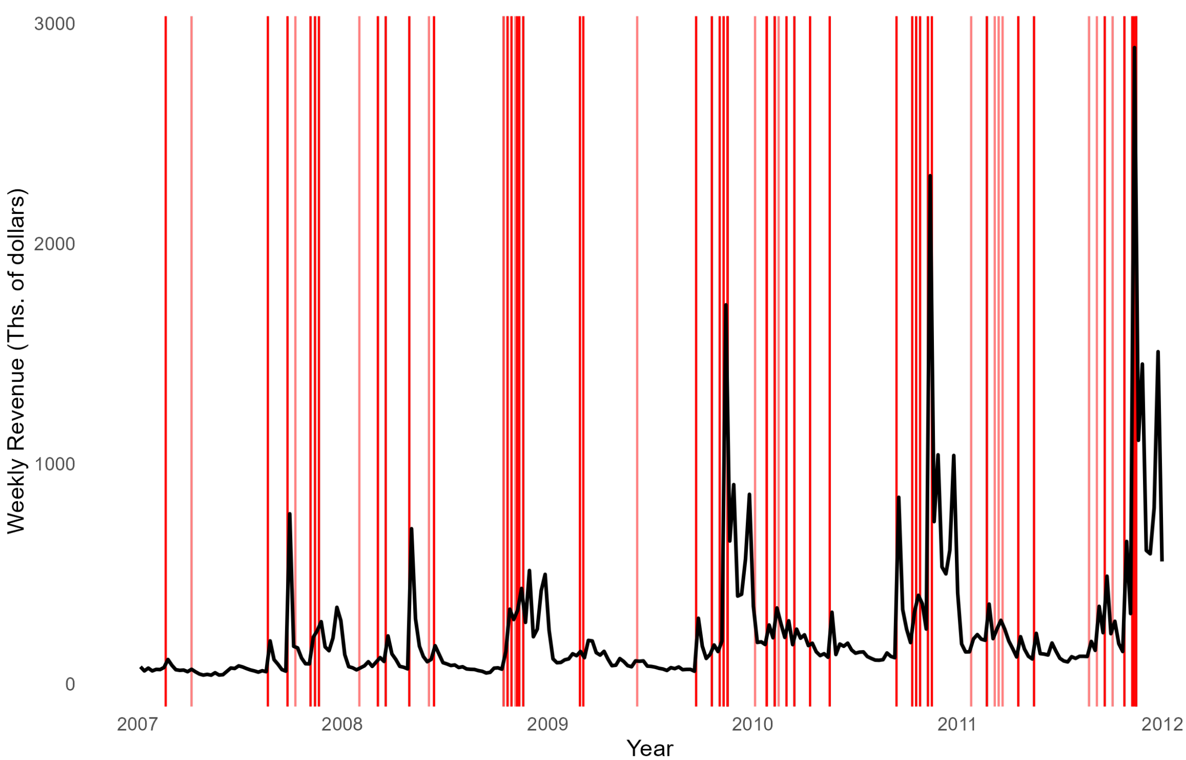

Figure 2: Graphical depiction of the first stage to predict new violent game sales

We instrument for the release month sale of new violent video games uses 2007 platform-specific software sales interacted with national game popularity, as measured in units sold. Our instrument is highly predictive of release month sales, even after adjusting for county-level video game spending and socioeconomic characteristics. In Panel B, we consider how the performance of our instrument changes if we use platform-specific software sales alone. This exercise reveals that patform-specific software sales alone are a predictive instrument for release month sales, but highlights that the interaction with national game popularity is useful for generating accurate predictions. See section 5 for details.

Panel A: Local platform popularity interacted with national game popularity

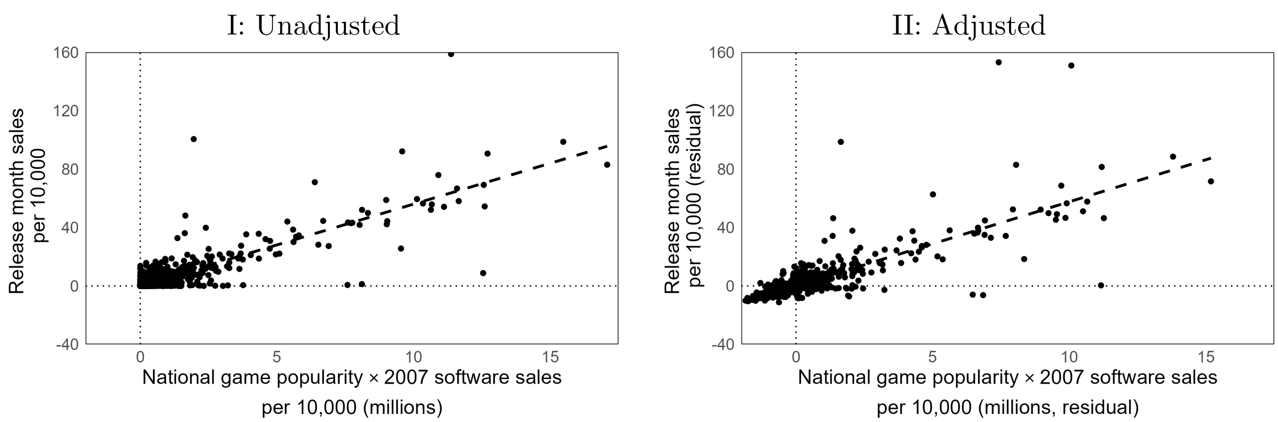

Panel B: Local platform popularity alone

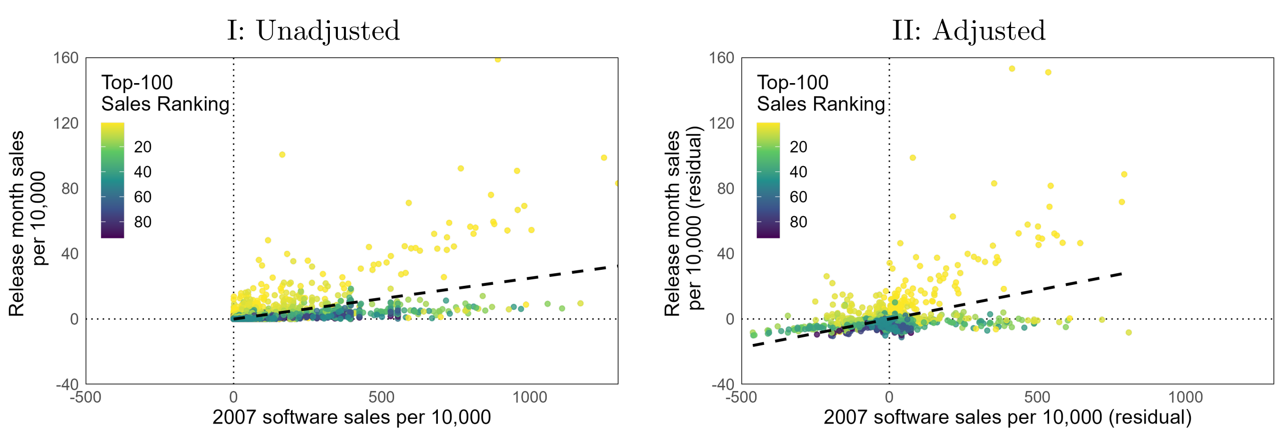

Figure 3: Variation in exposure to game consoles in 2007

There is systematic variation in exposure to new video games on the Xbox 360 and the PS3 as measured in Nielsen scanner data. As seen in Panel A, both new consoles have heavily concentrated sales in the Pacific Northwest. After controlling for the overall level of video game spending as well as county socioeconomic characteristics, there is no systematic pattern in exposure to game releases for either console. See section 5 for details.

Panel A: Unadjusted console game sales per 10,000 residents

I: Xbox 360

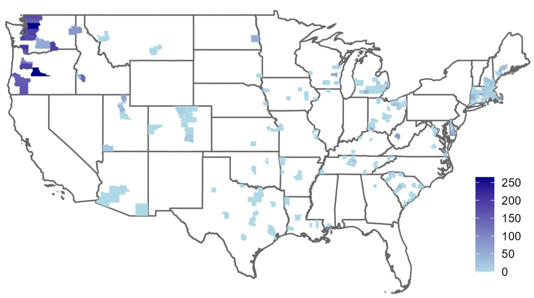

II: Playstation 3

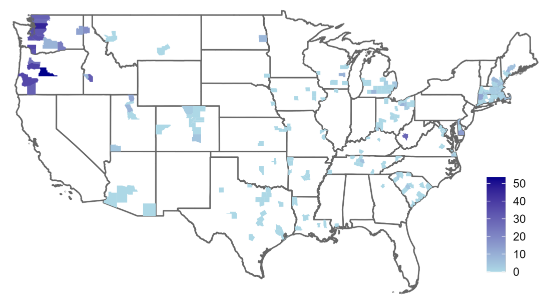

Panel B: Adjusted console game sales per 10,000 residents

I: Xbox 360

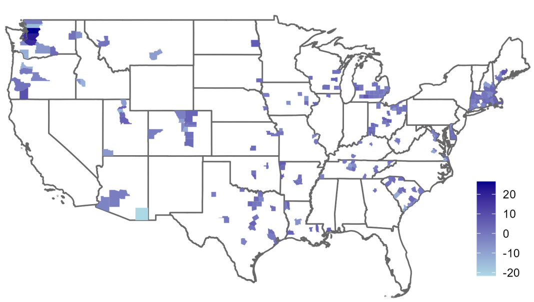

II: Playstation 3

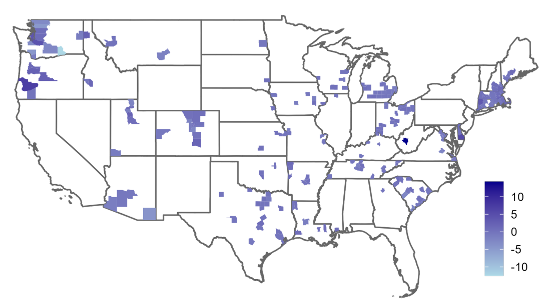

Figure 4: Effect size estimates of the effect violent game releases on violent crime, using 2007 platform shares

We estimate two-period difference-in-differences models to evaluate the effect of violent video game sales on violent crime rates. The proposed mechanisms through which violent games effect violent crime could vary in strength based on how much time has passed since game release. Thus we vary the length of the the pre- and post-period in the model, estimating the effect for periods of ± 8 weeks down to ± 1 week. Here, we have normalized the point estimates obtained from these models by the standard deviation in crime rates across agencies over corresponding period of time in 2007.. See Section 6 for details.

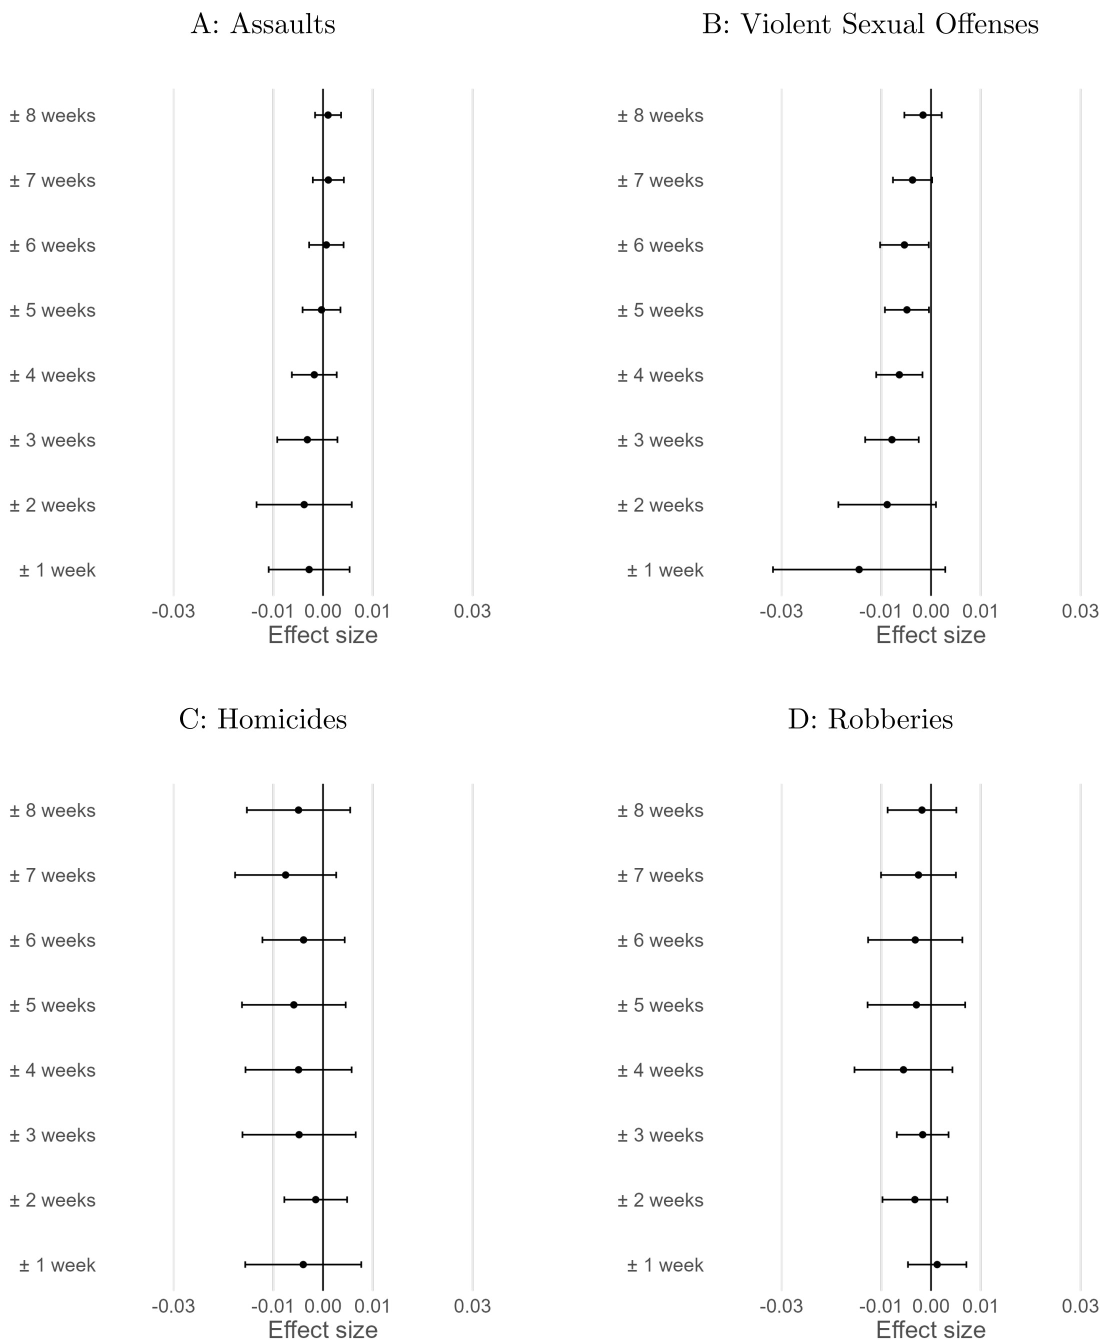

Figure 5: Impact estimates of the effect of violent game releases on violent crime, using 2007 platform shares

We estimate two-period difference-in-differences models to evaluate the impact of violent video game sales on violent crime rates. The proposed mechanisms through which violent games effect violent crime could vary in strength based on how much time has passed since game release. Thus we vary the length of the the pre- and post-period in the model, estimating the effect for periods of ± 8 weeks down to ± 1 week. To present our results in terms of impact, we have normalized the point estimates obtained from these models by the average crime rate across agencies over the corresponding period of time in 2007. See Section 6 for details.

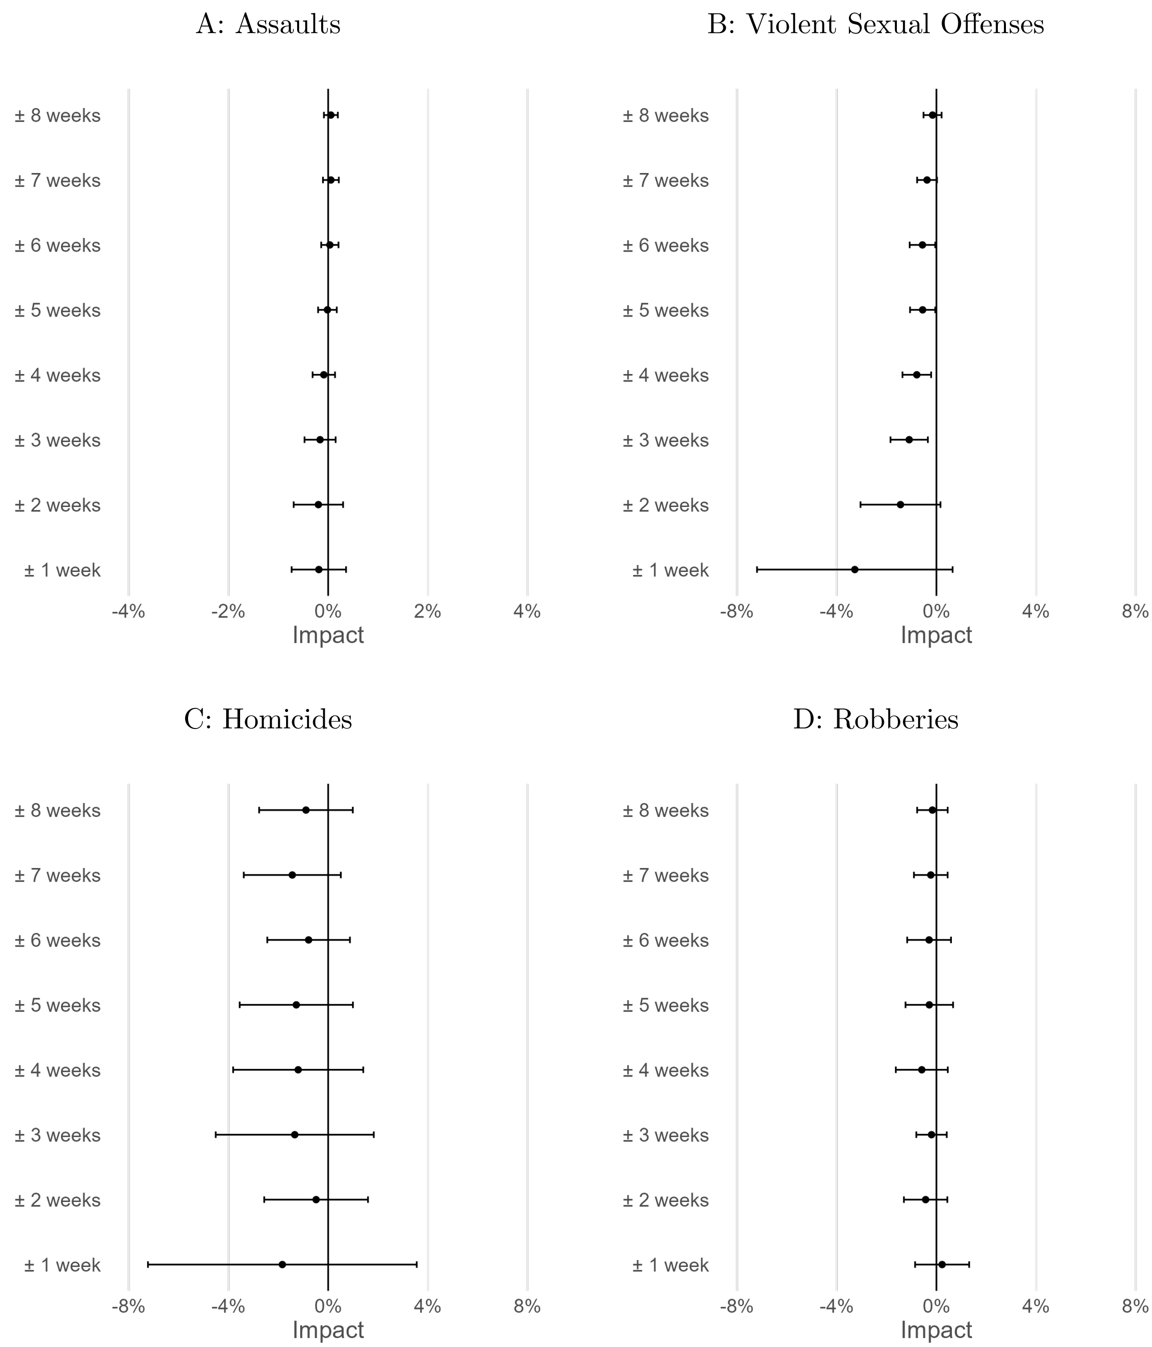

Figure 6: Effect size estimates of the effect violent game releases on violent crime, using 4-month lagged leave-one-out platform shares

We again estimate two-period difference-in-differences models to evaluate the effect of violent video game sales on violent crime rates, this time instrumenting for release sales using lagged platform-specific sales from the 16 weeks directly prior to each game release. Given that these lagged sales could be influenced by anticipatory behaviors, we leave out the 4 weeks directly prior to each game release. Here, we have normalized the point estimates obtained from these models by the standard deviation in crime rates across agencies over corresponding period of time in 2007. See Section 6 for details.

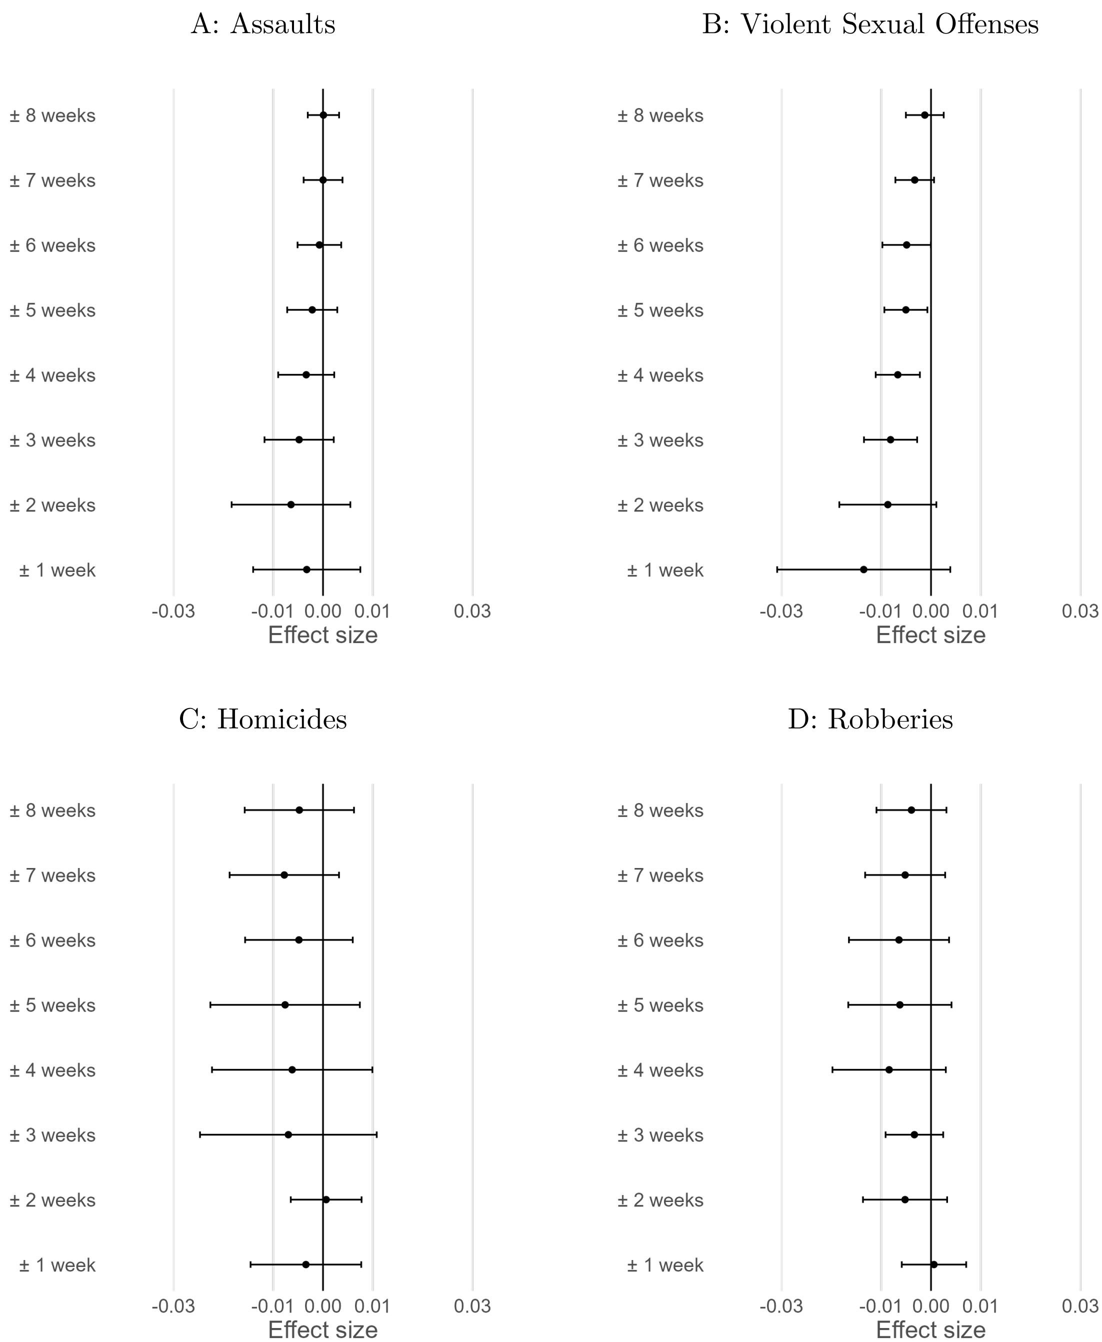

Figure 7: Impact estimates of the effect of violent game releases on violent crime, using 4-month lagged leave-one-out platform

We again estimate two-period difference-in-differences models to evaluate the effect of violent video game sales on violent crime rates, this time instrumenting for release sales using lagged platform-specific sales from the 16 weeks directly prior to each game release. Given that these lagged sales could be influenced by anticipatory behaviors, we leave out the 4 weeks directly prior to each game release. To present our results in terms of impact, we have normalized the point estimates obtained from these models by the average crime rate across agencies over the corresponding period of time in 2007. See Section 6 for details.

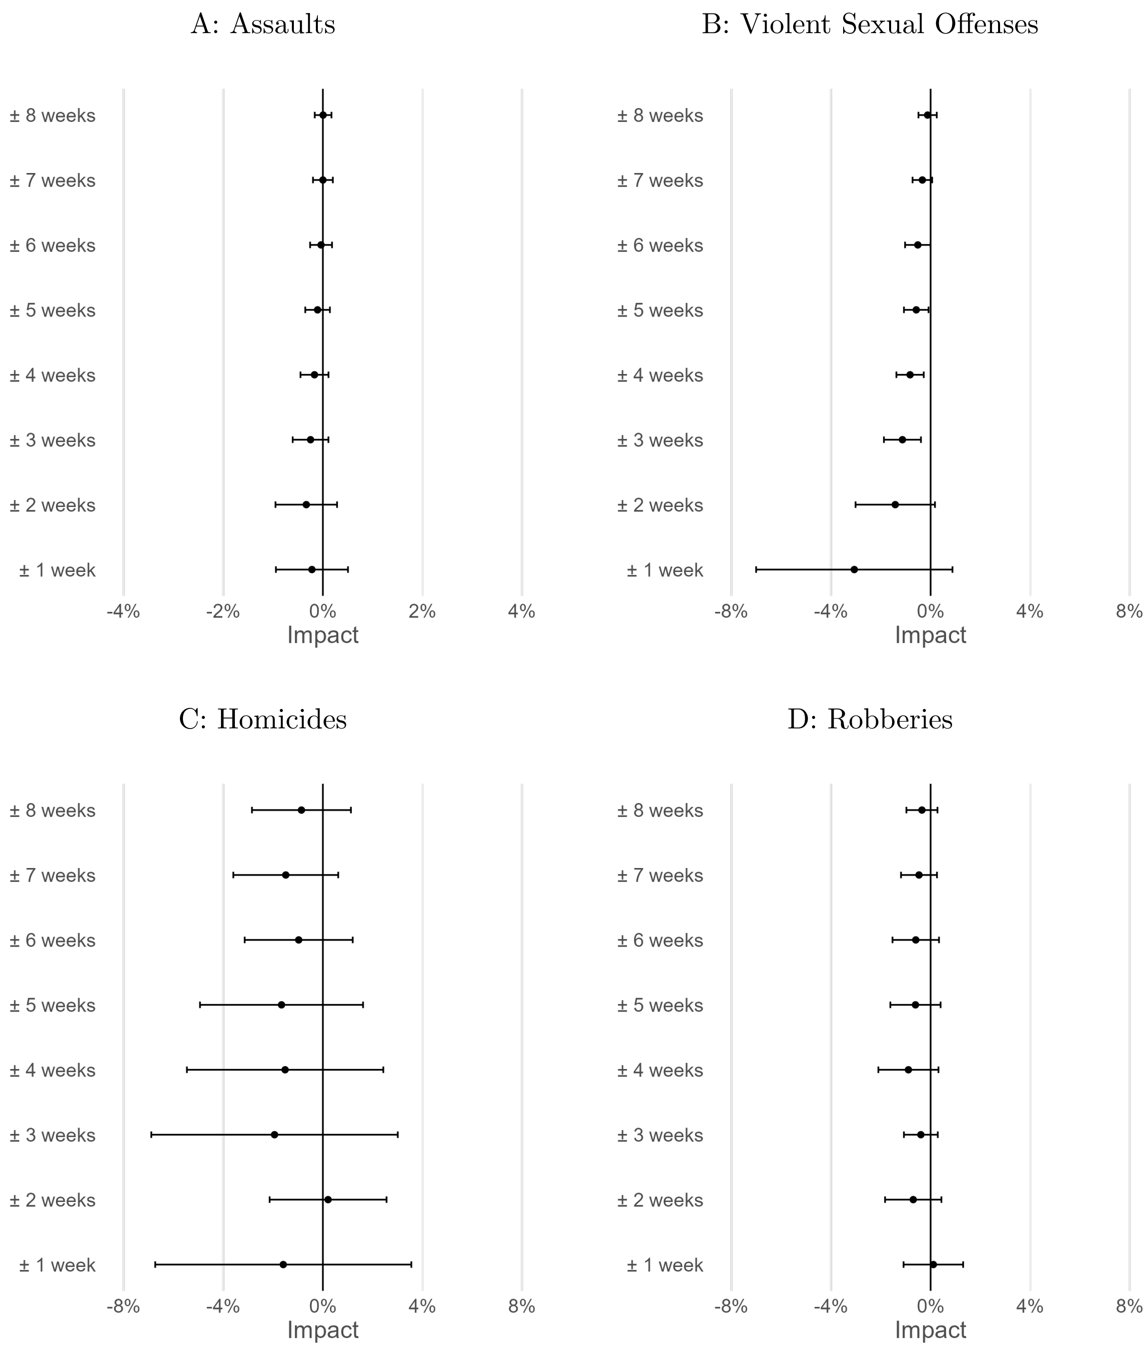

Figure 8: Event-study estimates of the effect of violent game releases on assault (using 2007 platform shares)

We use event study specifications to estimate the week-by-week effect of exposure to mature game releases. Most mature video games release in close proximity to other mature video games, complicating our ability to interpret point estimates. To avoid this challenge while retaining precision, we estimate our event study models using stratified samples for which there are different periods of time over which only one mature game releases. For instance, Panel A shows the results from the model using the 36 mature games that were released such that the next-closest mature game release was at least one week away. See Section 6 for details.

A: ±1 week (36 mature games)

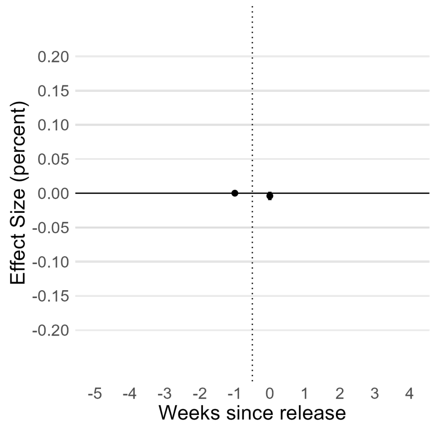

B: ±2 weeks (18 mature games)

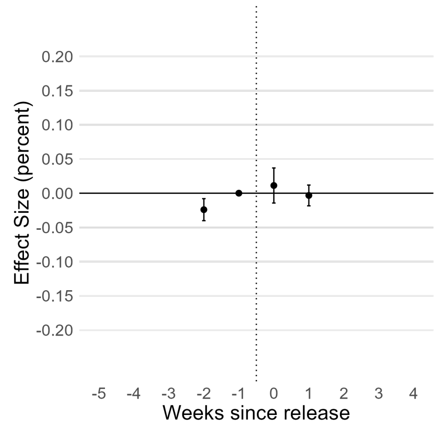

A: ±3 weeks (11 mature games)

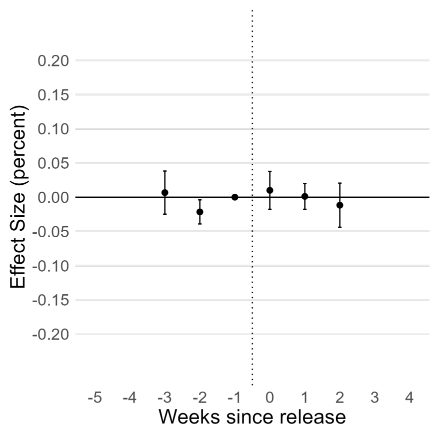

B: ±4 weeks (6 mature games)

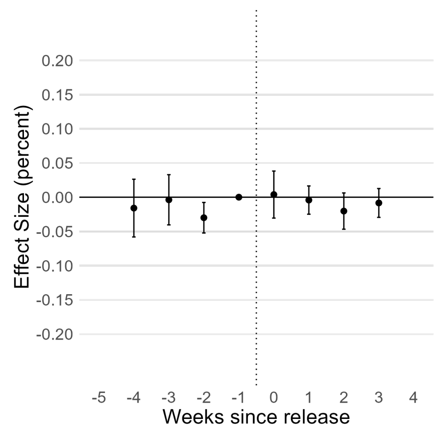

E: ±5 weeks (2 mature games)

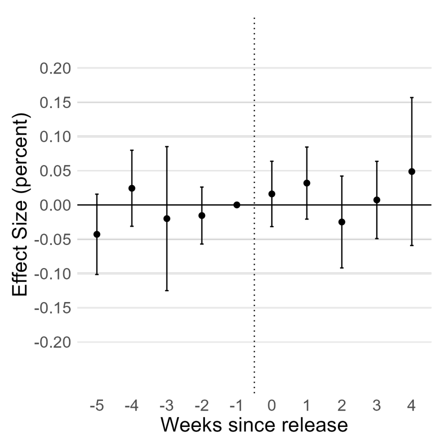

Figure 9: Event-study estimates of the effect of violent game releases on violent sex offenses (using 2007 video game sales as the instrument)

We use event study specifications to estimate the week-by-week effect of exposure to mature game releases. Most mature video games release in close proximity to other mature video games, complicating our ability to interpret point estimates. To avoid this challenge while retaining precision, we estimate our event study models using stratified samples for which there are different periods of time over which only one mature game releases. For instance, Panel A shows the results from the model using the 36 mature games that were released such that the next-closest mature game release was at least one week away. See Section 6 for details.

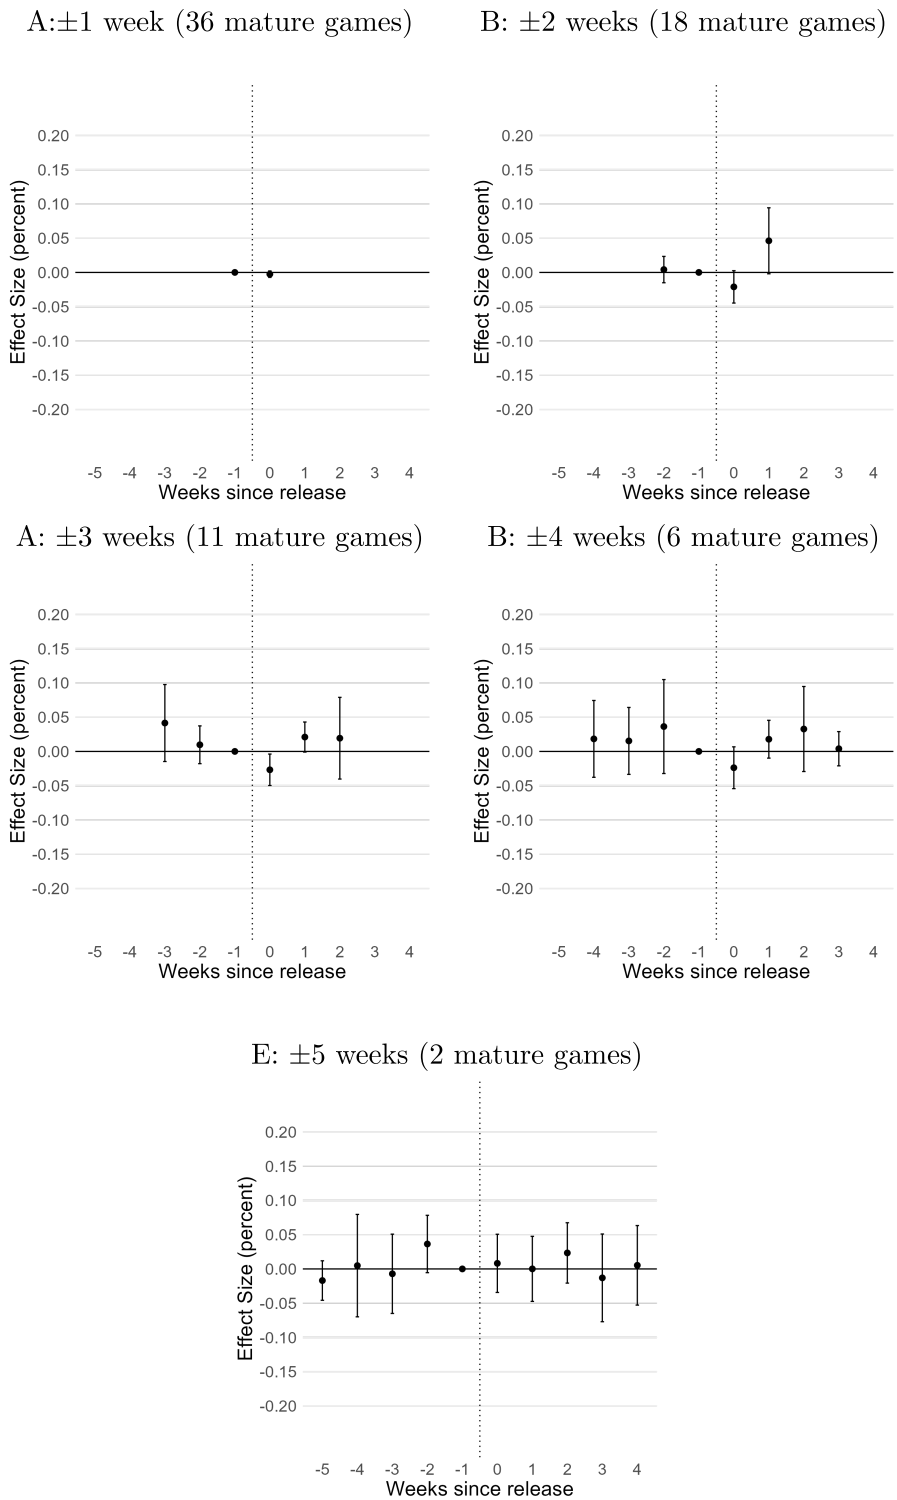

Figure 10: Event-study estimates of the effect of violent game releases on homicide (using 2007 platform shares)

We use event study specifications to estimate the week-by-week effect of exposure to mature game releases. Most mature video games release in close proximity to other mature video games, complicating our ability to interpret point estimates. To avoid this challenge while retaining precision, we estimate our event study models using stratified samples for which there are different periods of time over which only one mature game releases. For instance, Panel A shows the results from the model using the 36 mature games that were released such that the next-closest mature game release was at least one week away. See Section 6 for details.

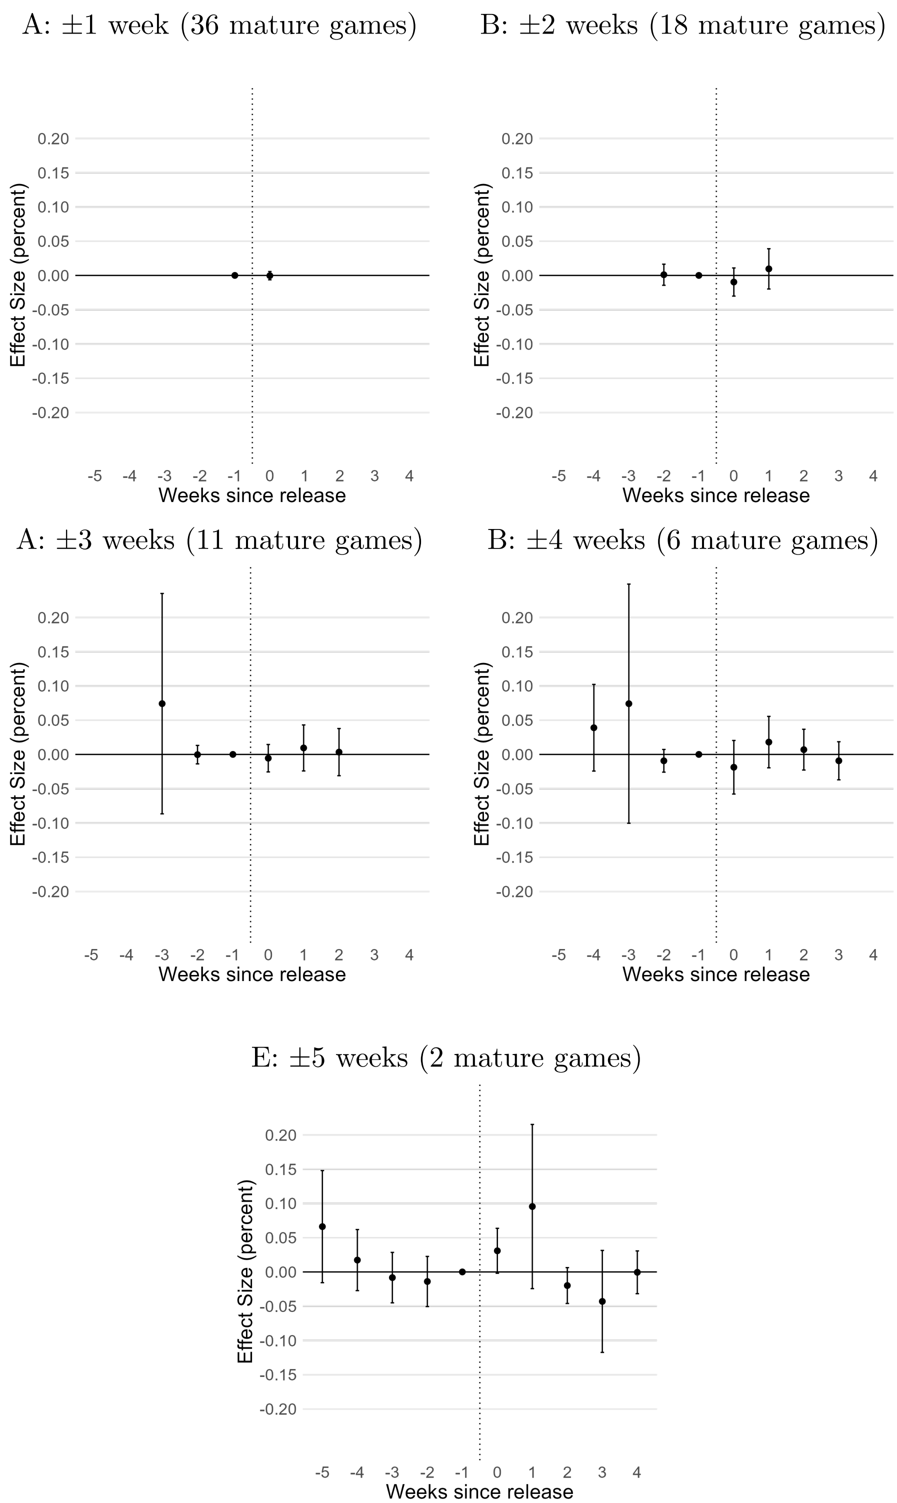

Figure 11: Event-study estimates of the effect of violent game releases on robberies (using 2007 video game sales as the instrument)

We use event study specifications to estimate the week-by-week effect of exposure to mature game releases. Most mature video games release in close proximity to other mature video games, complicating our ability to interpret point estimates. To avoid this challenge while retaining precision, we estimate our event study models using stratified samples for which there are different periods of time over which only one mature game releases. For instance, Panel A shows the results from the model using the 36 mature games that were released such that the next-closest mature game release was at least one week away. See Section 6 for details.

A: ±1 week (36 mature games)

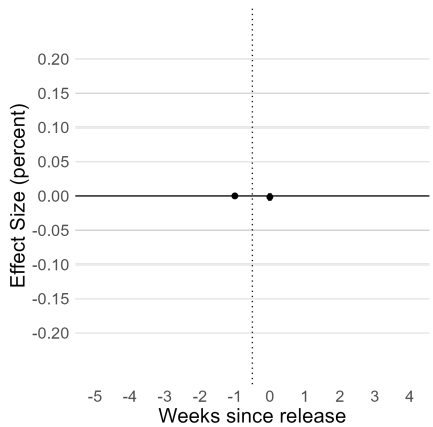

B: ±2weeks (18 mature games)

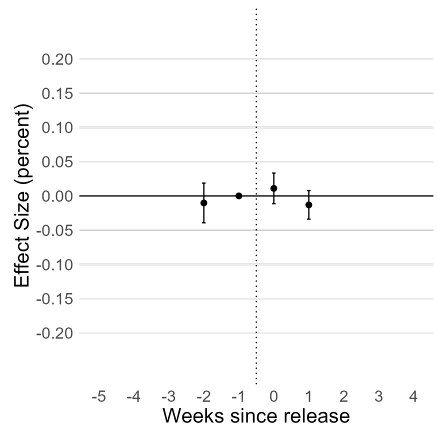

A: ±3 weeks (11 mature games)

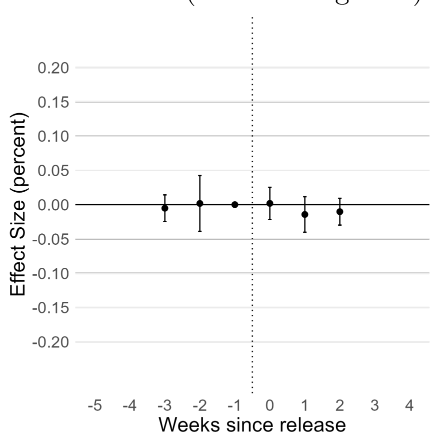

B: ±4 weeks (6 mature games)

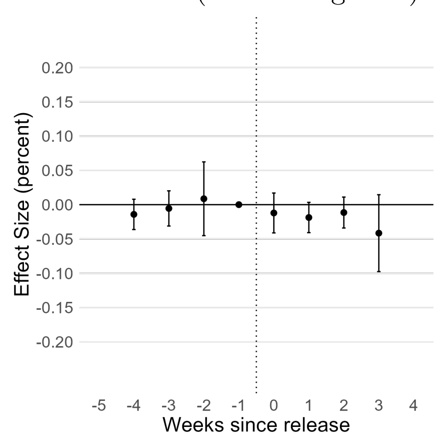

E: ±5 weeks (2 mature games)

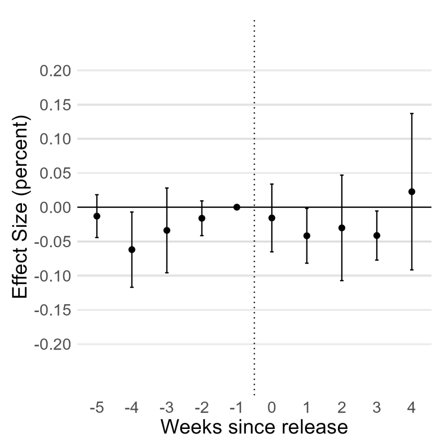

Table 1: Top-selling mature game releases 2007-2011

Name	Yearly Sales Rank	Release Date	Exclusive Release?
Crackdown	29	2007-02-20	Yes
Resident Evil 4: Wii Edition	66	2007-04-07	Yes
BioShock	32	2007-08-21	Yes
Halo 3	2	2007-09-25	Yes
The Orange Box	78	2007-10-09	No
Call of Duty 4: Modern Warfare	5	2007-11-05	No
Assassin's Creed	11	2007-11-13	No
Mass Effect	30	2007-11-20	No
Devil May Cry 4	70	2008-01-31	No
Army of Two	51	2008-03-04	No
God of War: Chains of Olympus	48	2008-03-04	Yes
Tom Clancy's Rainbow Six: Vegas 2	37	2008-03-18	No
Grand Theft Auto IV	6	2008-04-29	No
Ninja Gaiden II	91	2008-06-03	Yes
Metal Gear Solid 4: Guns of the Patriots	23	2008-06-12	Yes
SOCOM: U.S. Navy SEALs - Confrontation	95	2008-10-14	No
Saints Row 2	75	2008-10-14	No
Fable II	28	2008-10-21	Yes
Far Cry 2	96	2008-10-21	No
Fallout 3	38	2008-10-28	No
Resistance 2	65	2008-11-04	Yes
Gears of War 2	8	2008-11-07	Yes
Call of Duty: World at War	7	2008-11-11	No
Left 4 Dead	40	2008-11-18	No
Killzone 2	31	2009-02-27	Yes
Resident Evil 5	18	2009-03-05	No
Prototype	54	2009-06-09	No
Halo 3: ODST	8	2009-09-22	Yes
Borderlands	29	2009-10-20	No
Dragon Age: Origins	32	2009-11-03	No
Call of Duty: Modern Warfare 2	2	2009-11-10	No
Call of Duty: Modern Warfare Reflex Edition	85	2009-11-10	No
Assassin's Creed II	15	2009-11-17	No
God of War Collection	44	2009-11-17	Yes
Left 4 Dead 2	16	2009-11-17	No
Darksiders	88	2010-01-05	No
Heavy Rain	65	2010-01-25	No
Mass Effect 2	29	2010-01-26	Yes
BioShock 2	42	2010-02-09	No
Dante's Inferno	100	2010-02-09	No
Aliens vs. Predator	93	2010-02-16	No
Battlefield: Bad Company 2	22	2010-03-02	No
God of War III	15	2010-03-16	Yes
Tom Clancy's Splinter Cell: Conviction	49	2010-04-13	No
Red Dead Redemption	13	2010-05-18	No
Halo: Reach	5	2010-09-14	Yes
Medal of Honor	39	2010-10-12	No
Fallout: New Vegas	33	2010-10-19	No
Fable III	20	2010-10-26	Yes
Call of Duty: Black Ops	2	2010-11-09	No
Assassin's Creed: Brotherhood	20	2010-11-16	No
Dead Space 2	61	2011-01-25	No
Bulletstorm	94	2011-02-22	No
Killzone 3	37	2011-02-22	Yes
Dragon Age II	77	2011-03-08	No
Homefront	58	2011-03-15	No
Crysis 2	80	2011-03-22	No
Mortal Kombat	40	2011-04-19	No
SOCOM 4: U.S. Navy SEALs	99	2011-04-19	No
L.A. Noire	32	2011-05-17	No
Deus Ex: Human Revolution	87	2011-08-23	No
Dead Island	56	2011-09-06	No
Gears of War 3	6	2011-09-20	Yes
Rage	68	2011-10-04	No
Battlefield 3	7	2011-10-25	No
Call of Duty: Modern Warfare 3	1	2011-11-08	No
The Elder Scrolls V: Skyrim	8	2011-11-11	No
Assassin's Creed: Revelations	19	2011-11-15	No
Saints Row: The Third	53	2011-11-15	No
Halo: Combat Evolved Anniversary	44	2011-11-15	Yes

Notes: Games included and their yearly sales rank are based on top-selling games lists published by VGchartz.com. Staus as an exclusive release based on data scraped from gematsu.com.

Table 2: Top game releases by content rating between 2006 and 2011

ESRB Content rating	Content rating descripion	Share of top games with this content rating	Share present in sales data
E: for Everyone	*"Content is generally suitable for all ages. May contain minimal cartoon, fantasy or mild violence and/or infrequent use of mild language."*	54.6%	67.7%
T: for Teen	*"Content is generally suitable for ages 13 and up. May contain violence, suggestive themes, crude humor, minimal blood, simulated gambling and/or infrequent use of strong language."*	19.3%	89.1%
M: for Mature	*"Content is generally suitable for ages 17 and up. May contain intense violence, blood and gore, sexual content and/or strong language."*	26.1%	97.3%

Notes: Set of top games determined using yearly top-100 lists from VGchartz.com. Content rating descriptions are determined and published by the Electronic Software Rating Board (ESRB). Games with the E-10 ESRB rating are grouped with E games. See Section 2.1 for more details on the ESRB system.

Table 3: First-stage estimates for video game sales during release month

	(1)	(2)	(3)
Panel A: 2007 instrument			
National game popularity · 2007 software sales	0.0000057***	0.0000056***	0.0000056***
	(0.0000004)	(0.0000004)	(0.0000004)
Per-10,000 game spending		0.0000116	0.0000130
		(0.0000083)	(0.0000082)
SES Controls	No	No	Yes
F-stat	22188	16290	16314
N	9693	9693	9693
Panel B: Lagged instrument			
National game popularity· lagged software sales	0.0000209***	0.0000216***	0.0000216***
	(0.0000009)	(0.0000011)	(0.0000011)
Per-10,000 game spending		-0.0001082***	-0.0001078***
		(0.0000242)	(0.0000243)
SES Controls	No	No	Yes
F-stat	29143	27049	27004
N	9838	9838	9838

Notes: Observations here are at the county level to be consistent with our source of identifying variation (video game sales across counties). Socioeconomic (SES) controls are per capita personal income, the poverty rate, and the unemployment rate. Standard-error estimates allow for clusters at the county level. The reported F-statistic is for the exclusion of the shift-share instrument interacting 2007 software sales with national game popularity. *, **, and ***, indicate statistical significance at the ten-, five-, and one-percent levels, respectively