Journal of Monetary Economics 63 (2014) 1–18

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Securitization and moral hazard: Evidence from credit score cutoff rules

Ryan Bubb [a,*], Alex Kaufman [b,1]

[a] New York University, 40 Washington Square South, New York, NY 10012, United States
[b] Board of Governors of the Federal Reserve System, 20th Street and Constitution Avenue N.W., Washington, D.C. 20551, United States

ARTICLE INFO

Article history:
Received 9 March 2012
Received in revised form
22 January 2014
Accepted 23 January 2014
Available online 8 February 2014

Keywords:
Financial crisis
Moral hazard
Mortgages
Securitization
Credit scores

ABSTRACT

A growing literature exploits credit score cutoff rules as a natural experiment to estimate the moral hazard effect of securitization on lender screening. However, these cutoff rules can be traced to underwriting guidelines for *originators*, not for securitizers. Moreover, loan-level data reveal that lenders change their screening at credit score cutoffs in the absence of changes in the probability of securitization. Credit score cutoff rules thus cannot be used to learn about the moral hazard effect of securitization on underwriting. By showing that this evidence has been misinterpreted, our analysis should move beliefs away from the conclusion that securitization led to lax screening.

1. Introduction

It has become conventional wisdom that securitization contributed to the sharp rise in mortgage defaults that precipitated the recent financial crisis. The logic of the moral hazard problem posed by securitization seems straightforward: lenders that sell loans they originate to dispersed investors may bear less of the cost when loans default and hence may have less incentive to screen borrowers. The belief that this moral hazard problem played an important role in the financial crisis has influenced regulatory reform, with the 2010 Dodd-Frank Act adopting a requirement that securitizers retain a 5% interest in mortgages they securitize to better align their incentives.[2]

However, securitization may not have had a large moral hazard effect on underwriting. Economists usually believe that moral hazard causes profitable trade to not occur, or that it leads to the development of incentive mechanisms to overcome the problem. And indeed, mortgage lenders and securitizers developed a range of practices to mitigate moral hazard (Gorton, 2009), including contractual provisions as well as software systems that automate mortgage underwriting, and achieved a high level of trade years before the crisis without apparent incident.[3]

* Corresponding author. Tel.: +1 212 9928871.
E-mail addresses: ryan.bubb@nyu.edu (R. Bubb), alex.kaufman@gmail.com (A. Kaufman).
[1] Tel.: +1 617 7774851.
[2] Section 941 of the Dodd-Frank Wall Street Reform and Consumer Protection Act, codified at 15 U.S.C. Section 78o-11. This law includes a set of exceptions to this requirement and also applies this requirement to other types of securitizations.
[3] As far back as 1993, nearly two-thirds (65.3%) of mortgage volume was packaged into MBS, about the same fraction as in 2006 (67.6%) on the eve of the crisis (2010 *Mortgage Market Statistical Annual*, published by Inside Mortgage Finance). The contractual structures of securitization did become more complex over time, but the same potential moral hazard problem posed by securitization existed decades ago.

http://dx.doi.org/10.1016/j.jmoneco.2014.01.005
0304-3932

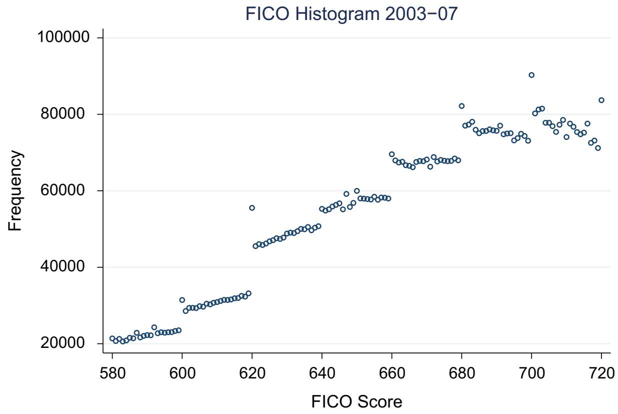

Fig. 1. Discontinuities in the density of mortgages by FICO credit score. Data source: Lender Processing Services Applied Analytics, Inc. Sample of first-lien, non-buydown, owner-occupied, single-family mortgage loans originated between January 2003 and December 2007.

Moreover, there are other plausible explanations for the loose lending of the pre-crisis years. For instance, lenders, securitizers, and investors alike may have been over-optimistic about housing prices, causing them to underestimate the risk of low-downpayment loans (Foote et al., 2012). Hence it is an empirical question whether moral hazard due to securitization, as opposed to some other factor, was the key driver of the decline in underwriting standards leading up to the crisis.

This paper investigates the most influential evidence to-date on the moral hazard effect of securitization, which is based on discontinuities in lender behavior at certain credit scores. Credit scores are used by lenders as a summary measure of default risk, with higher credit scores indicating lower default risk. Despite the smoothness of the distribution of credit scores in the overall population, histograms of mortgage borrower credit scores, such as Fig. 1, exhibit a series of large discontinuities.

A number of recent papers employ a regression discontinuity (RD) research design that exploits these discontinuities as a "natural experiment" to learn about the moral hazard effect of securitization on lender screening (Keys et al., 2009; Krainer and Laderman, 2009; Keys et al., 2010b; Jiang et al., 2010; Rajan et al., 2010; Keys et al., 2012). This research design is based on a particular theory for the origin of these discontinuities—the *securitization rule-of-thumb theory*. First offered by Keys et al. (2010b), this theory posits that private-label securitizers employ a rule of thumb whereby they are exogenously more willing to purchase loans made to borrowers with FICO[4] scores just above 620 than those to borrowers with scores just below 620. Based on this theory, these papers interpret jumps in mortgage default at credit score thresholds as establishing that securitization led to moral hazard in screening. This literature has been highly influential among economists and policymakers. For example, Treasury Secretary Timothy Geithner cited this literature in a recent report supporting the Dodd-Frank Act's new securitization risk retention requirements.[5]

However, crucial to the validity of this RD design is the assumption that lenders' discontinuous change in screening at these credit score cutoffs is exclusively driven by a change in the probability of securitization at the cutoff. If there is another reason for lenders' behavior to change discontinuously, then the jump in defaults at the cutoff cannot be attributed to securitization. In the terminology of instrumental variables (IV) this assumption is the exclusion restriction.

We develop an alternative theory for the origin of credit score cutoff rules—the *origination rule-of-thumb theory*. Institutional evidence shows that in the 1990s, with the goal of improving underwriting, Fannie Mae and Freddie Mac (the Government-Sponsored Enterprises, or GSEs) required originators to adopt credit score cutoff rules to determine how carefully to screen mortgage borrowers. These cutoff rules were later incorporated into widely used automated underwriting software, and in time they became industry-wide origination standards.

A simple model based on discreteness in the cost of information collection provides a simple explanation for why the GSEs directed originators to adopt such credit score cutoff rules. Furthermore, the discontinuity in lender screening at the credit score cutoffs creates discontinuities in the amount of private information originators have about loans. Information asymmetry can inhibit trade, and the model shows that origination rules of thumb can thus result in discontinuities in the securitization rate.

The origination rule-of-thumb theory and the institutional evidence on which it is based show that the evidence for moral hazard based on credit score cutoff rules has been misinterpreted. The independent change in lender screening intensity at the cutoffs violates the exclusion restriction, invalidating an RD design based on these cutoffs. The jumps in

[4] The credit scoring model developed by Fair Isaac and Company (FICO) is the industry standard.

[5] Timothy Geithner, "Macroeconomic Effects of Risk Retention Requirements," January 2011. The report notes that "subprime borrowers with credit scores just above a threshold commonly used by securitizers to determine which loans to purchase defaulted at significantly higher rates than those with credit scores below the threshold" (p. 11). The report concludes "that markets are unable, in certain circumstances, to align the incentives of parties in the securitization chain adequately" and that "such weaknesses demonstrate the need for regulatory reforms" (p. 14).

default at credit score thresholds thus do not provide evidence that securitization led to lax screening. By showing that the most influential evidence that securitization led to moral hazard is premised on a crucial assumption that does not hold, our analysis should move beliefs away from the conclusion that mortgage securitization caused poor underwriting through a moral hazard channel.

The origination rule-of-thumb theory is further corroborated by quantitative evidence. The paper that first used the securitization rule-of-thumb theory (Keys et al., 2010b) used a dataset of only securitized loans and hence could not show whether there were discontinuities in the securitization rate at the credit score cutoffs used by lenders. Loan-level data from Lender Processing Services (LPS) that includes both securitized loans and non-securitized "portfolio" loans exhibit large jumps in the number and default rate of loans at credit score cutoffs in the absence of corresponding discontinuities in the securitization rate. This is strong evidence that lenders are using origination rules of thumb that are unrelated to any change in the probability of securitization.

In IV terms, the first stage shows that the instrument (the credit score cutoff) has no effect on the treatment (securitization). Despite this, it has a large reduced-form effect on the outcome (lender screening), confirming the exclusion restriction violation implied by the institutional evidence. These results are inconsistent with the securitization rule-of-thumb theory, but consistent with the origination rule-of-thumb theory.

The progenitors of the securitization rule-of-thumb theory wrote a response to an earlier draft of this paper (Keys et al., 2012) in which they argue that securitization rules of thumb were only used by private-label securitizers, and that a particular subgroup of loans in the LPS dataset—low documentation loans originated with the intent to sell to private-label securitizers—reveals discontinuities in securitization and lender screening, in line with their theory.

However, identifying a subgroup of loans for which there *is* a discontinuity in the probability of securitization would not make the RD design valid for that subgroup. Key to the RD research design is the assumption that if there were no moral hazard effect on lender screening then there would be no jumps in default at credit score cutoffs. In fact, there are discontinuities in lender screening at 620 and 660 in the absence of discontinuities in securitization for a wide variety of samples, including the jumbo market, which is exclusively a private securitization market, and for loans originated in 2008–2009 after the collapse of the private securitization market. This failure of the exclusion restriction means that the RD research design cannot provide evidence on the causal hypothesis of interest. Furthermore, as explained in more detail in Section 4.4 below, Keys et al. (2012)'s selection of only low documentation loans entails selecting on the outcome of interest—lender screening—and hence is econometrically incorrect. Moreover, their methodology for isolating a pure private-label securitization market is based on the assumption that no loans are at risk of being securitized both by private-label securitizers and by the GSEs, and this assumption is inconsistent with the evidence.

The paper proceeds as follows. Section 2 first describes the securitization rule-of-thumb theory and the RD research design based on it. Section 3 then describes our origination rule-of-thumb theory and the institutional evidence on which it is based and shows that it implies a violation of the exclusion restriction of this RD research design. Section 4 presents quantitative evidence using a large sample of mortgage loans that shows the existence of discontinuities in lender screening in the absence of discontinuities in the securitization rate, confirming the violation of the exclusion restriction implied by the institutional evidence. Responses to issues raised in papers written in response to this paper, as well as additional details, are provided in the online appendix. Section 5 concludes.

2. Securitization rule-of-thumb theory

Why did lenders adopt credit score cutoff rules? The theory currently most accepted in the finance literature, first offered in Keys et al. (2010b), is that lenders adopted cutoff rules solely in response to a rule of thumb followed by securitizers in their purchasing decisions. Under this theory, the rule of thumb made securitizers more willing to buy mortgage loans made to borrowers with FICO scores of 620 or above than loans to those with scores of 619 or below, and the higher probability of selling loans above 620 reduced lenders' incentive to screen those loans.

This rule of thumb started with the guidelines of the GSEs. As Keys et al. (2010b, pp. 318–319) explains:

> Guidelines by Freddie Mac suggest that FICO scores below 620 are placed in the *Cautious Review Category*, and Freddie Mac considers a score below 620 "as a strong indication that the borrower's credit reputation is not acceptable." … We argue that adherence to this cutoff by subprime MBS investors, following the advice of GSEs, generates an increase in demand for securitized loans that are just above the credit cutoff relative to loans below this cutoff. There is widespread evidence that is consistent with 620 being a rule of thumb in the securitized subprime lending market. … The credit threshold of 620 was used by nearly all the [top fifty subprime] lenders."

We refer to this theory as the *securitization rule-of-thumb theory*. Importantly, the securitization rule-of-thumb theory assumes that lenders did not respond to these guidelines directly. Rather, only private-label securitizers and MBS investors followed this advice of the GSEs by adopting the cutoffs as purchasing rules of thumb. In turn, subprime lenders faced a discontinuous increase in the probability that they could sell loans above 620, which through a moral hazard effect resulted in subprime lenders adopting a 620 cutoff rule in screening applicants.

If the securitization rule-of-thumb theory were true, then one could use the resulting variation in secondary market demand for loans as a natural experiment to estimate the causal effect of securitization on lender screening. The pools of potential borrowers with FICO scores of 619 and 620 are very similar. But when originators consider applicants from each of these pools, the probability they can sell a 620 loan is higher than the probability they can sell a 619 loan because of the rule of thumb followed by purchasers in the secondary market. The moral hazard effect of securitization can then be estimated by comparing the performance of 619 loans to 620 loans. And indeed, Keys et al. (2010b) and others show that defaults jump up at 620 and interpret this as demonstrating that securitization had a moral hazard effect on lender screening.

However, this research design is based on a crucial identifying assumption: there must be no reason for the default rate to change discontinuously at the 620 threshold other than in response to a change in ease of securitization. If instead there is a reason for the default rate to change discontinuously at 620 even if the securitizers' behavior did not change at that cutoff, then one cannot estimate the causal effect of any change in probability of securitization at 620 by estimating the change in default rates at 620.

This assumption can be expressed in IV terms as an exclusion restriction. The instrument is an indicator for whether a loan is made to a borrower with a FICO score greater than or equal to 620. The endogenous treatment of interest is probability of securitization. Probability of default is the ultimate outcome of interest. The exclusion restriction necessary for this instrument to be valid is that the 620 cutoff does not have an effect on defaults through any channel other than its effect on the probability of securitization.

The securitization rule-of-thumb theory predicts that one should observe both discontinuities in the securitization rate and, if there is a moral hazard effect of securitization, also in lender screening at 620 FICO. The securitization rule-of-thumb theory can be rejected if there are discontinuities in lender screening at the credit score cutoffs established by the GSEs in the absence of a discontinuity in the securitization rate, as this would show that lender screening changes at the cutoffs for reasons other than a change in the probability of securitization.

Importantly, no paper that relies on the securitization rule-of-thumb theory as the basis of its identification strategy has provided institutional evidence that private securitizers adopted a purchasing rule-of-thumb in response to the GSEs' guidelines. Keys et al. (2010b, pp. 318–319), for instance, only cites evidence that *originators*, not private securitizers, followed credit score cutoff rules in making lending decisions. Moreover, they used a dataset of only securitized loans and thus could not show whether there is a discontinuity in the securitization rate.

3. Origination rule-of-thumb theory

We articulate an alternative *origination rule-of-thumb theory* for the origin of credit score cutoff rules based on the institutional evidence: lenders responded directly to the underwriting guidelines of the GSEs, which required lenders to adopt credit score cutoff rules to determine how carefully to screen borrowers, and over time those cutoff rules became industry-wide origination standards. Importantly, this theory and the institutional evidence on which it is based imply that originators change their screening behavior at credit score cutoffs for reasons other than a change in the probability of securitization. This means that the exclusion restriction of the RD research design described in Section 2 is invalid and therefore that credit score cutoff rules cannot be used to estimate the causal effect of securitization on lender screening.

3.1. Institutional evidence on the origin of credit score cutoff rules

The credit score cutoff rules used by lenders have their origin in underwriting guidelines for lenders established by the GSEs. In the mid-1990s, prior to the adoption of credit scoring in mortgage underwriting, the GSEs conducted research into the relationship between FICO scores and mortgage performance showing that "despite the fact that those borrowers who had FICO scores in the lower range (620 or less) represented only a very small percentage of the total universe, they (as a group) accounted for approximately 50% of the eventual defaults…" (Fannie Mae, 1995, p. 4).

In 1995 Freddie Mac sent a letter to originators directing them to begin using credit scores in underwriting and establishing three tiers of credit scores. The key part of the letter is provided in the online appendix. The FICO scores of 620 and 660 were important cutoffs. For borrowers with FICO scores above 660, lenders were to do a "basic" review to "underwrite the file as required to confirm the borrower's willingness to repay as agreed." For borrowers with FICO scores between 660 and 620, lenders were to perform a "comprehensive" review to "underwrite all aspects of the borrower's credit history to establish the borrower's willingness to repay as agreed." For borrowers with FICO scores below 620, lenders were warned to be "cautious" and to "perform a particularly detailed review of all aspects of the borrower's credit history to ensure that you have satisfactorily established the borrower's willingness to repay as agreed." Fannie Mae (1997, pp. 8–9) established a similar set of cutoffs, including at both 620 and 660 FICO. Lenders who sold loans to Fannie Mae and Freddie Mac were contractually obligated to follow the GSEs' guidance letters establishing credit score cutoff rules for screening.

Credit score cutoff rules spread in large part through their incorporation into automated underwriting systems (AUSs), which became widely adopted beginning in the mid-1990s (Hutto and Lederman, 2003). Most mortgage originators use either the Desktop Underwriter (DU) program, created by Fannie Mae, or the Loan Prospector (LP) program, created by Freddie Mac. These and other similar programs take as inputs information such as FICO score, loan-to-value ratio, and debt-to-income ratio, and compute a recommendation. When lenders get an "approve" or "accept" recommendation from their AUS that is usually the end of the process and they approve the loan. When they receive a "refer" or "caution" recommendation, they may then begin the process of manual underwriting (Hutto and Lederman, 2003).

Because DU and LP were designed and distributed by the GSEs, which require lenders to use 620 and 660 as cutoffs, these and other credit score cutoffs were coded directly into the AUS decision rules.[6] Though AUSs calculate default risk using smooth functions of FICO score, they also employ a layer of "overwrites" which trigger a "refer" recommendation when borrowers fall into certain categories—for instance, borrowers with FICO scores below 620. As a result, lenders following AUS recommendations are discontinuously more likely to initiate manual underwriting for a borrower with 619 FICO than for a borrower with 620 FICO. Loans that are "referred" by the AUS are still eligible for purchase by the GSEs (and private securitizers) so long as the lender judges them to be acceptable through its manual underwriting process. Securitizers commonly buy loans that are initially referred and later approved through the manual underwriting process.

3.2. A rational model of credit score cutoff rules

Why did the management of the GSEs believe that using credit score cutoff rules to determine how carefully to screen potential borrowers would help originators "better assess and manage the quality of [their] loan originations" (Freddie Mac, 1995, p. 2)? In general, lenders face nondivisible, discrete screening decisions: whether to manually underwrite a loan, whether to conduct a face-to-face interview, and so on. We posit that the discreteness of such decisions implies that the optimal lender screening strategy takes the form of a credit score cutoff rule or set of cutoff rules.

This rationale for credit score cutoff rules can be formalized in a simple model. Suppose that there is a continuum of prospective borrowers of unit mass. Each borrower has a credit score x that represents hard information about the likelihood that the borrower will default. To economize on notation, let $x \in [0, 1]$ represent both the borrower's credit score and the fraction of borrowers with that credit score who would repay a mortgage. For simplicity, the remaining $1-x$ fraction of borrowers with that credit score would default without making any payments on the mortgage. x is distributed according to the strictly positive, continuous probability density function $f(x)$. Borrowers would like to take out a mortgage for 1 unit of the numeraire good at time 0 to be repaid with interest at time 1, but they have an outside option such that they will refuse a loan offer with a gross interest rate above $\overline{R} > 1$. There is a single risk-neutral lender with discount factor normalized to 1. At time 0 each borrower applies to the lender for a mortgage. The lender observes each applicant's x.

The lender then chooses whether to further investigate each borrower's creditworthiness. To do so, the lender must bear a fixed cost $c > 0$ per applicant. The fixed cost arises from discreteness in the information production function available to the firm managers who set underwriting policy. For example, requiring loan officers to meet with loan applicants in person, or to perform manual underwriting in addition to the commonly used computer-aided automated underwriting process, entails a fixed cost per applicant. Moreover, it would be difficult for managers (or the GSEs) to specify continuous investigation intensities for continuous distributions of borrowers, given difficulty in monitoring their agents' screening behavior (Ellison and Holden, 2008). Consequently, firm managers face a discrete choice set of investigation intensities.

If the lender investigates and the borrower is a defaulter, the lender learns this with probability $s \in (0, 1)$, and otherwise the lender observes nothing. The lender's investigation thus reveals this "defaulter signal" about a borrower with credit score x with probability $(1-x)s$. Assume that $c < (\overline{R}-1)s/\overline{R}$ so that investigation is cheap enough that it will pay for the lender to investigate some applicants. The lender then chooses whether to lend to each applicant and, if so, makes a take-it-or-leave-it interest rate offer $R(x)$. Those offered loans then decide whether to accept the offer. In period 1, borrowers learn whether they are defaulters, and the nondefaulters pay the lender $R(x)$.

Obviously the lender never chooses to lend to applicants for which its investigation revealed the defaulter signal. Furthermore, because the lender has all of the bargaining power, it should be obvious that, if the lender lends, it is a dominant strategy to offer $\overline{R}$, and for all borrowers offered a loan to accept. Hence, the equilibria of the game are characterized by an investigation strategy (which borrower types the lender investigates) and a lending strategy (to which types the lender offers loans). Our main result is

Proposition 1. *In the unique equilibrium, the lender uses cutoff rules based on a lending threshold* $\underline{x} = (1-s+c)/(\overline{R}-s)$ *and a screening threshold* $\overline{x} = 1-c/s > \underline{x}$:

1. *The lender rejects borrowers with* $x < \underline{x}$.
2. *The lender investigates borrowers with* $\underline{x} \le x < \overline{x}$ *and offers loans to those for which its investigation does not reveal the defaulter signal.*
3. *The lender offers loans to borrowers with* $x \ge \overline{x}$ *without investigation.*

All proofs are in the online appendix.

This screening behavior by lenders results in a discontinuous jump in the density of loans at the $\overline{x}$ screening threshold and a similar jump in the default rate of loans. The intuition for how these discrete costs result in discontinuities in default rates is straightforward: if lenders gave stricter scrutiny to loan applicants just above the $\overline{x}$ threshold it would reduce the

[6] Personal communication with Freddie Mac executives, October 9th, 2009. Fannie Mae changed DU in 2000 to calculate its own risk rating based on the underlying credit report data (Quinn, 2000), but FICO scores are still required to evaluate loans in DU (Fannie Mae, 2007).

default rate, but this reduction would not justify bearing the fixed cost c per applicant to collect the information. In contrast, for loan applicants just below the $\overline{x}$ threshold the benefit of additional information outweighs the fixed cost.

Our simple model provides a rationale for the GSEs' expressed belief that requiring originators to follow a credit score cutoff rule would improve underwriting. As with the credit score x in the model, there is a continuous, monotonic relationship between FICO score and default risk. Mapped into our model, a FICO score such as 620 corresponds to the screening threshold $\overline{x}$.

Consider now the implications of origination rules-of-thumb for securitization. The lender's greater investigation of applicants below the screening threshold results in the lender having greater private information about those loans than the loans above the cutoff. As in Akerlof (1970), asymmetry of information on borrower quality can inhibit trade. And because the amount of private information changes discontinuously at the screening threshold, this adverse selection problem can result in the volume of trade—that is, the securitization rate—to change discontinuously at the same threshold.

A similar result obtains if the lender and securitizer bargain over a loan purchase contract prior to the lender making investigation and lending decisions, creating a potential moral hazard problem. This can be formalized in a simple extension to our model. Suppose a securitizer exists with a cost of funds slightly less than the lender's cost of funds, so that its discount factor is $\delta = 1+\varepsilon$ for arbitrarily small ε. The securitizer and lender bargain over a contract characterized by two functions and an up-front payment: $\sigma(x)$ denotes the fraction of loans of type x that the securitizer will purchase, $T(x)$ represents the price that it will pay, and T represents an up-front payment that determines the ultimate division of surplus between the securitizer and lender. The game then proceeds as in the baseline model but, after loans are made, the lender sells a fraction $\sigma(x)$ of loans of each type x to the securitizer for a payment $T(x)$ per loan, with the securitizer choosing the particular loans that it purchases randomly at each x.

First suppose that the lender's screening behavior is contractible. In this case, then the securitizer and lender will implement the first-best screening behavior characterized above, and the securitizer will buy all of the lender's loans, as stated formally in the following proposition.

Proposition 2. *In the equilibrium of the model with a securitizer and contractible lender screening, the lender's behavior is the same as in the model without securitization, given in Proposition 1, and the fraction of loans securitized is $\sigma(x) = 1$ for all $x > \underline{x}$.*

Suppose now that there is asymmetric information so that lender screening behavior is not contractible. Securitization now raises a moral hazard problem for lender screening, and the securitizer will leave a fraction of loans on the lender's books to maintain its incentives to screen, as formalized in the following proposition.

Proposition 3. *In the equilibrium of the model with a securitizer and non-contractible lender screening, the lender's behavior is the same as in the model without securitization, given in Proposition 1, and the fraction of loans securitized for each x is given by*

$$\sigma^*(x) = \begin{cases} \dfrac{\overline{R}s(1-x)x - c}{\overline{R}s(1-x)x} & \text{if } \underline{x} \le x < \overline{x} \\ 1 & \text{if } x \ge \overline{x} \end{cases}$$

Importantly, the fraction of loans securitized jumps discontinuously at the screening threshold $\overline{x}$. When it is efficient for the lender to extend a loan without investigation (that is, $x \ge \overline{x}$), there is no moral hazard problem, and the securitizer purchases all of the loans. When it is efficient for the lender to investigate (that is, $\underline{x} \le x < \overline{x}$), the securitizer purchases a fraction of loans for each value of x such that the remaining portfolio loans provide sufficient incentive for the lender to investigate.

The origination rule-of-thumb theory thus predicts that there will be discontinuities in lender screening at the credit score cutoffs established by the GSEs. The theory is also consistent with the existence or not of discontinuities in the securitization rate at those same credit score thresholds, depending on whether asymmetric information between the lender and the securitizer inhibits trade in mortgages.

3.3. Discussion

The origination rule-of-thumb theory and the institutional evidence on which it is based imply that one cannot use credit score cutoffs to learn about the moral hazard effect of securitization. The GSEs' guidelines establishing the credit score cutoff rules were not directed at investors or private securitizers; rather, they were screening guidelines for *originators*.[7] Moreover, these credit score cutoff rules were incorporated into the GSEs' contracts with originators and also encoded into

[7] Note the tension between Keys et al.'s view that the 620 purchasing rule of thumb was only used by private-label securitizers and not the GSEs, and Keys et al.'s institutional story that the rule-of-thumb originated with advice from the GSEs. To believe the securitization rule-of-thumb theory, you have to believe that (1) the GSEs' direction to lenders to adopt the 620 rule-of-thumb had no direct effect on lender behavior, (2) the GSEs themselves did not follow a *purchasing* rule-of-thumb at 620, and (3) private-label securitizers did listen to the 620 guidance that the GSEs gave to lenders about screening and adopted a rule-of-thumb in purchasing at 620. This chain of events is implausible, even without considering the institutional and quantitative evidence against it presented in this paper.

underwriting software that was widely used in the industry, including for loans sold to private securitizers. In time these credit score cutoffs became industry-wide origination standards. This means that lender behavior changes at the credit score cutoffs for reasons other than any change in probability of securitization, which violates the exclusion restriction of the RD research design described in Section 2 above. Therefore, contra a large and influential literature in finance, the jumps in default at certain credit score cutoffs do not demonstrate that securitization caused moral hazard in underwriting in the run-up to the financial crisis.

The fact that these cutoff rules began with guidance from the GSEs, which are themselves securitizers, does not imply that cutoff rules can nonetheless be used in an RD design to learn about the moral hazard effect of securitization on lender screening. These credit score cutoff rules do not provide any exogenous variation in probability of securitization that can be used to test the moral hazard hypothesis. Rather, the cutoff rule evidence shows that the GSEs were, to some extent, able to control originators' underwriting behavior. When Fannie Mae and Freddie Mac concluded that originators could improve their underwriting practices by adopting credit score cutoff rules, they used a range of mechanisms to get originators to do so. The evidence shows that the GSEs' efforts resulted in an industry-wide adoption of the credit score cutoff rules. Rather than providing evidence for moral hazard, the cutoff rules provide evidence that the GSEs, which bore the credit risk on the loans they securitized, were to some extent successful in their efforts to require lenders to use specific underwriting methodologies that the GSEs believed would improve loan origination.

4. Quantitative evidence

The institutional evidence on the origin of credit score cutoff rules alone provides strong evidence against the securitization rule-of-thumb theory and the exclusion restriction of the RD research design based on it. Quantitative evidence can be used to further test the competing theories. Both theories for the origin of credit score cutoff rules predict that there are discontinuities in lender screening at credit score thresholds. Furthermore, discontinuities in the securitization rate are consistent with both theories. However, the securitization rule-of-thumb theory places a further restriction on the data: it predicts that wherever there is a discontinuity in lender screening there will also be a discontinuity in securitization. Therefore, if there are discontinuities in lender screening in the absence of a discontinuity in securitization, the data would reject the securitization rule-of-thumb theory. Such a pattern, however, would be consistent with the origination rule-of-thumb theory and the institutional evidence. The evidence in this section using a large loan-level dataset shows that this pattern appears all over the mortgage market.

4.1. Empirical strategy

Our empirical strategy is summarized below.

4.1.1. Data

Our data come from Lender Processing Services Applied Analytics, Inc. (LPS). These are loan-level data collected from participating mortgage servicers, including the 10 largest servicers in the United States. LPS contains privately securitized loans, GSE-securitized loans, and portfolio loans (loans for which the originator retains rights to the payment stream).[8] The data cover over half of outstanding mortgages in the United States and contain more than 32 million active loans. Key variables in the dataset include borrower FICO scores, detailed loan terms, monthly securitization status, and monthly loan performance data.

Our sample is all first-lien, non-buydown, owner-occupied, single-family mortgage loans originated between January 2003 and June 2009 in the LPS dataset. Borrowers must have FICO scores between 500 and 800 to be included in the sample.[9] In order to avoid survivorship bias, loans that do not appear in the dataset by their sixth month after origination are omitted from the sample, even if they appear later.[10]

Our main sample is loans originated between January 2003 and December 2007. This pre-crisis period is of obvious importance for understanding the causes of the mortgage crisis. The post-crisis period is also useful. Because private-label mortgage securitization shut down in 2008, the set of loans originated between January 2008 and June 2009 is used to investigate whether the use of credit score cutoff rules disappeared along with private-label securitization. Table 1 provides summary statistics for our data. Additional analysis is provided in the online appendix showing that our sample from LPS is roughly representative of the overall market.

[8] The data does not distinguish between loans that GSEs hold in their own portfolio and loans that the GSEs securitize. However, the distinction is irrelevant for our purposes. For simplicity the term "securitized" refers to refer to all loans purchased on the secondary market.

[9] 14.8% of loans from LPS that meet our other sample selection criteria have missing FICO scores and are dropped from our analysis sample.

[10] Loans that first enter the dataset significantly after they are originated (so-called "seasoned" loans) can potentially lead to biased estimates. Those loans least at risk of terminating early, either through repayment or foreclosure, are often not typical of the greater population of loans, and such loans will be over-represented in seasoned data. For this reason our sample includes only loans with seasoning of 6 months or fewer.

Table 1
Summary statistics: pre- and post-crisis samples.

	2003–2007			2008–2009H1		
	Mean	*S.D.*	*N*	Mean	*S.D.*	*N*
GSE securitized	0.609	0.488	14,498,143	0.946	0.227	3,437,910
Private securitized	0.284	0.451	14,498,143	0.013	0.115	3,437,910
Portfolio	0.107	0.310	14,498,143	0.041	0.198	3,437,910
Low-doc	0.304	0.460	9,193,406	0.228	0.419	2,723,308
Adjustable	0.292	0.455	14,498,143	0.036	0.185	3,435,707
Refi	0.546	0.498	14,498,143	0.589	0.492	3,437,910
FHA/VA	0.069	0.254	14,498,143	0.327	0.469	3,437,910
Borrower FICO	699.4	65.4	14,498,143	716.0	61.7	3,437,910
Loan amount ($)	246,869	217,381	14,498,060	212,685	143,725	3,437,908
Loan-to-value	74.1	17.3	14,338,067	77.3	19.1	3,388,371
Delinquent	0.130	0.336	14,498,143	0.046	0.210	3,437,910

Notes: Sample includes first-lien, non-buydown, owner-occupied, single-family mortgage loans. *GSE securitized*, *Private securitized*, and *Portfolio* capture securitization status 6 months after origination. *Low-doc* includes both "low" and "no" documentation loans. *Adjustable* refers to adjustable-rate mortgages. *Refi* refers to mortgage loans used to refinance an existing mortgage. *FHA/VA* refers to mortgages insured by the Federal Housing Administration or the Department of Veterans Affairs. *Borrower FICO* refers to the FICO credit score of the borrower on the loan. *Loan amount* in 2009 dollars. *Loan-to-value* is the ratio of the loan amount to the appraised value of the home securing the loan. For the 2003–2007 sample, *Delinquent* equals 1 if loan became 61 days or more overdue within 36 months of origination. For the 2008–2009 sample, because not enough time has yet elapsed, *Delinquent* uses a shorter follow-up period of 18 months after origination.

4.1.2. Measuring lender screening

Our main outcome of interest is lender screening intensity. Lenders using laxer standards are more likely to approve applicants, which leads to a higher lending rate per potential applicant in the population. Jumps in the lending rate create jumps in the total number of loans originated. Therefore our first measure of lender screening intensity is the total number of loans.

Our second measure of screening intensity is the default rate. This measure is appealing because filtering out applicants who will not repay is the ultimate goal of screening. Our definition of default is whether payment was delinquent by 61 days or more at any time during a loan's follow-up period. A loan's follow-up period extends from the month the loan was originated until 36 months after origination for our main 2003–2007 sample. For our January 2008–June 2009 sample, the follow-up period is only 18 months due to data limitations.[11]

4.1.3. Measuring the probability of securitization

Our measure of whether a loan is securitized is a binary variable equal to 1 if a loan is sold out of portfolio at any time during the loan's follow-up period and equal to 0 otherwise.

The probability a loan is securitized is the relevant determinant of lenders' incentives to use to investigate the moral hazard effects of securitization. In contrast, Keys et al. (2010b, 2012) argue that the right measure of probability of securitization is the "unconditional probability" of securitization (that is, the probability that a *potential borrower* is given a loan which is later securitized, rather than either not being given a loan at all or being given a loan that is kept in portfolio).[12] However, the unconditional probability conflates two different probabilities: (1) the probability that potential borrowers are given a loan (the lending rate); and (2) the probability that loans are securitized (the securitization rate). More formally, let $L_i \in \{0,1\}$ denote whether potential borrower i is given a loan and let $S_i \in \{0,1,\emptyset\}$ denote whether borrower i's loan is securitized (with $S_i = \emptyset$ if borrower i is not given a loan). The unconditional probability is then

$$Pr(S_i = 1) = Pr(L_i = 1) * Pr(S_i = 1 | L_i = 1) \quad (1)$$

The first factor on the right-hand side of this equation is the lending rate; the second factor is the securitization rate.

The probability relevant for testing the hypothesis that securitization weakened the screening incentives of lenders is the probability that a loan is securitized ($Pr(S_i = 1 | L_i = 1)$), not the probability a potential borrower is given a securitized loan ($Pr(S_i = 1)$). If a lender has a very high probability of selling a loan, say to a naive securitizer, then the lender's incentives to screen borrowers might be attenuated. If instead there is a large chance that the lender will keep the loan, then the moral hazard problem is less severe. In equilibrium, what matters for lenders' incentives is the probability they have to keep the loan.

[11] Our dataset was obtained in early 2011, when December 2010 was the most recent available month.

[12] A natural measure of the "unconditional probability" of securitization at each FICO score would be the number of securitized loans at each FICO score divided by the number of potential borrowers at each FICO score. However, Keys et al. (2010b, 2012) simply use the number of securitized loans at each FICO score as a proxy for the "unconditional probability" of securitization, arguing that the number of potential borrowers is continuous in FICO score and thus any discontinuous jump in the number of securitized loans must therefore reflect a jump in the "unconditional probability" of securitization.

4.1.4. Specifications

Our statistical test for discontinuities in the number of loans at credit score thresholds is based on McCrary (2008), which develops a formal test of the continuity of the density function of the running variable in RD analyses that allows for proper inference.

To estimate discontinuities in the probability of default and the probability of securitization, a standard RD analysis using OLS is performed (i.e., the linear probability model). For the 620 FICO cutoff, 6th-order polynomials on either side of the cutoff are estimated using the full sample:

$$Y_i = \beta_0 + \beta_1 \mathbb{1}_{\{FICO_i \geq 620\}} + f(FICO_i) + \mathbb{1}_{\{FICO_i \geq 620\}} * g(FICO_i) + \epsilon_i \tag{2}$$

where i indexes individual loans, the dependent variable Y_i indicates either whether loan i defaulted or whether it was securitized, and both $f(FICO_i)$ and $g(FICO_i)$ are 6th-order polynomials in *FICO*. An analogous specifications is used to estimate discontinuities at the FICO score of 660.[13]

4.2. Main results

Table 2 presents estimates using our main sample of loans originated in the pre-crisis 2003–2007 period. There are large jumps in the number of loans at both 620 and 660, indicating increases in originators' lending rate. There are also sharp discontinuities in the default rate at the cutoffs. The jump at 620, for instance, is approximately 3.5 percentage points on a base of 23.5 percentage points—a nearly 15% increase. Together these show that lenders dramatically change their screening behavior at 620 and 660.

However, there is no evidence of a jump up in the securitization rate. Indeed, the estimates of securitization rate discontinuities at 620 and 660 are both economically small and *negative*. For example, below 620 lenders have a 92.3% chance of selling the loans they make, whereas above 620 lenders have a 91.6% chance of selling the loans. Fig. 2 graphically illustrates our finding of discontinuities in lending and in default, but not in the probability of securitization.

These findings corroborate the institutional evidence that shows that lender screening changes at the 620 and 660 cutoffs for reasons other than a change in the probability of securitization. Relatedly, this is direct evidence of the failure of the exclusion restriction assumption necessary for the securitizer rule-of-thumb theory to be valid. Despite there being no “first stage” discontinuity in the probability of securitization, there is a large reduced-form effect of credit score cutoffs on lender screening. This is because the cutoffs (the instrument) have an effect on lender screening (the outcome) through channels other than probability of securitization (the treatment).

4.3. Private-label vs. GSE subgroups

Keys et al. (2010a, 2012), written in response to an earlier draft of this paper, argue that the GSEs and the private-label securitizers operated in completely separate markets, and that only private-label securitizers, and not the GSEs, followed a 620 rule of thumb in making purchasing decisions. The authors argue that our finding of jumps in the number of loans and the default rate in the absence of a corresponding jump in the securitization rate is an artifact of inappropriately pooling loans bought by private-label securitizers with loans bought by the GSEs.

4.3.1. Breakdown by realized securitization status

To investigate this, we begin by simply breaking our sample into three subgroups based on securitization status: private-label securitized loans, GSE-purchased loans, and loans kept in the portfolio of the originating lender.[14] Table 3 presents estimates of the lending discontinuity and the default discontinuity for these three subgroups. Note that securitization rates cannot be calculated for these subgroups because they have been selected on securitization status. Fig. 3 presents the results graphically.

The default discontinuity point estimate at 620 for the private-label subgroup (3.3 percentage points) is slightly larger than the point estimate for the GSE subgroup (2.2 percentage points), though viewed as an increase over the base rate of default the private-label subgroup has a slightly smaller increase (10.2%) than the GSE subgroup (11.8%).[15] The largest increase, in both absolute and percent terms, is in the portfolio subgroup. There is a 6.5 percentage point, and 26.4%, increase in default at 620 among loans that are never securitized. Results for FICO 660 show statistically indistinguishable percentage-point jumps in all three subgroups, with the GSE group having the largest jump in percent terms. All subgroups exhibit large jumps in the number of loans at both 620 and 660.

[13] Results are similar using a local linear regression in which the sample is restricted to a 10 FICO score point band on either side of the threshold (unreported for brevity).

[14] Identity of securitizer is measured at 6 months after origination. All results are robust to alternate definitions.

[15] Keys et al. (2012) also investigated whether GSE loans exhibit a discontinuity at 620 FICO and found no discontinuity in a sample from LPS originated in 2001–2006. However, we are unable to replicate this finding, even using a 2001–2006 sample meant to mimic Keys et al. (2012)'s sample selection. We are unaware of any reason for this discrepancy in results. We asked the authors for the code used to produce their results, but they did not provide it.

Table 2
Discontinuities in main 2003–2007 pre-crisis sample.

	FICO 620			FICO 660		
	log (# *loans*) (1)	*Delinquent* (2)	*Securitization* (3)	log (# *loans*) (4)	*Delinquent* (5)	*Securitization* (6)
Change at cutoff	0.322***	0.035***	−0.007***	0.122***	0.024***	−0.004***
s.e.	(0.003)	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)
Base rate	–	0.235	0.923	–	0.161	0.922
N	14,498,143			14,498,143		

Notes: Sample is loans originated in January 2003–December 2007. *Delinquent* is an indicator for whether the loan ever became 61 days delinquent within 36 months of origination. Securitization is an indicator for whether the loan was ever sold within 36 months of origination. Columns 1 and 4 use a local linear regression to test for a discontinuity in the number of loans at the credit score cutoff, as outlined in McCrary (2008). Columns 2, 3, 5, and 6 fit a 6th-order polynomial in FICO on either side of the cutoff. "Base rate" refers to the level of the dependent variable right below the credit score cutoff. Heteroskedasticity-robust standard errors in parentheses.
* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

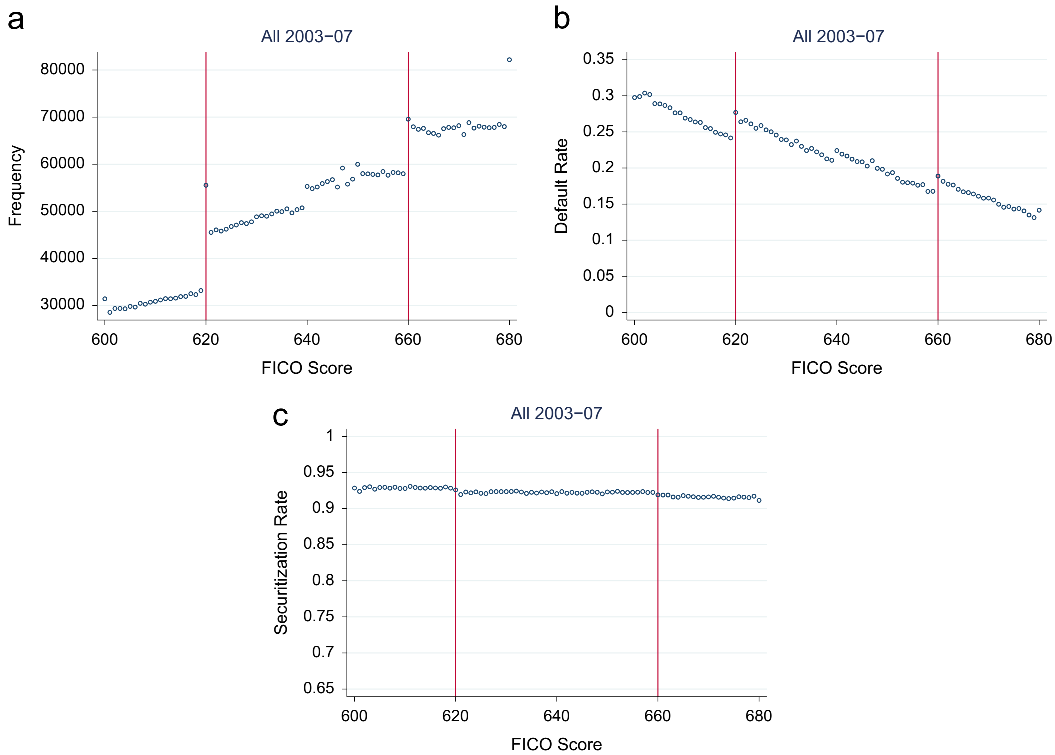

Fig. 2. Main sample. Data source: Lender Processing Services Applied Analytics, Inc. Sample of first-lien, non-buydown, owner-occupied, single-family, mortgage loans originated between January 2003 and December 2007. A loan is considered defaulted if it ever became 61 days delinquent within 36 months of origination. A loan is considered securitized if it was ever sold within 36 months of origination: (a) Frequency. (b) Default. (c) Securitization.

These results make clear that discontinuities in lender screening at credit score cutoffs are not solely driven by loans that become private-label securitized. To the contrary, these jumps in the number of loans and default rate are a robust feature of all loans, including those securitized by the GSEs and those kept in portfolio.

Additionally, these figures provide a good illustration of why the securitization rate, rather than the "unconditional probability of securitization," is the correct measure of securitization. Imagine for a moment one were investigating the opposite hypothesis from Keys et al. (2010b), namely that a *lack* of securitization—that is, lenders having to keep loans in portfolio—causes moral hazard. Keys et al. (2010b) uses the total number of securitized loans as its measure of the ease of securitization. Applying this same logic to our opposing hypothesis, one would use the total number of portfolio loans as our measure of how difficult it is to securitize. If there were jumps in the number of portfolio loans and in their default rate at credit score cutoffs, one would analogously interpret this as evidence that a lack of securitization destroys lenders' incentives to screen.

Table 3
Discontinuities in subgroups: private-label, GSE, and portfolio.

	FICO 620		FICO 660	
	log (*# loans*) (1)	*Delinquent* (2)	log (*# loans*) (3)	*Delinquent* (4)
Panel A: Private-Label (N=4,124,533)				
Change at cutoff	0.370***	0.033***	0.181***	0.019***
s.e.	(0.005)	(0.004)	(0.004)	(0.003)
Base rate	–	0.324	–	0.249
Panel B: GSE (N=8,824,642)				
Change at cutoff	0.288***	0.022***	0.077***	0.018***
s.e.	(0.004)	(0.003)	(0.003)	(0.001)
Base rate	–	0.186	–	0.115
Panel C: Portfolio (N=1,548,968)				
Change at cutoff	0.342***	0.065***	0.181***	0.020***
s.e.	(0.008)	(0.007)	(0.008)	(0.004)
Base rate	–	0.246	–	0.194

Notes: Sample is loans originated in January 2003–December 2007. The sample used in Panel A includes only loans that were private-label securitized. The sample used in Panel B includes only loans that were purchased and/or securitized by the GSEs. The sample used in Panel C includes only loans that were kept by the originator. Securitization status is measured 6 months after origination. *Delinquent* is an indicator for whether the loan ever became 61 days delinquent within 36 months of origination. Columns 1 and 3 use a local linear regression to test for a discontinuity in the number of loans at the credit score cutoff, as outlined in McCrary (2008). Columns 2 and 4 fit a 6th-order polynomial in FICO on either side of the cutoff. "Base rate" refers to the level of the dependent variable right below the credit score cutoff. Heteroskedasticity-robust standard errors in parentheses.
* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

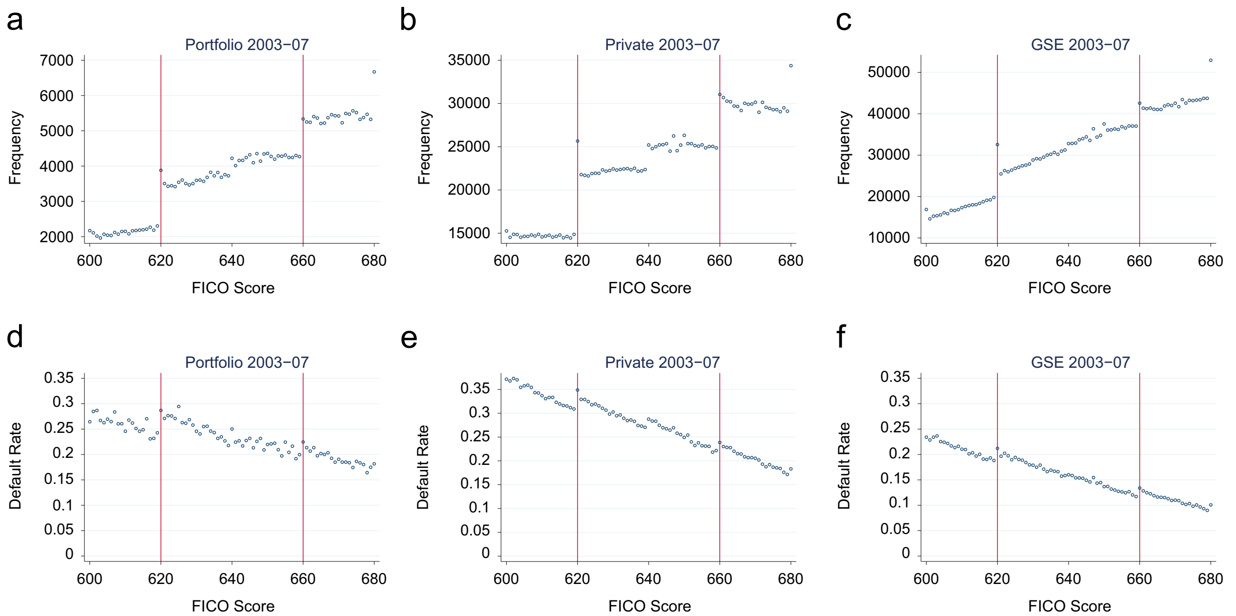

Fig. 3. Frequency and default discontinuities for portfolio, private-label securitized, and GSE securitized loans originated in 2003–2007. Data source: Lender Processing Services Applied Analytics, Inc. Sample of first-lien, non-buydown, owner-occupied, single-family mortgage loans originated between January 2003 and December 2007. A loan is considered defaulted if it ever became 61 days delinquent within 36 months of origination. A loan is considered securitized if it was ever sold within 36 months of origination: (a) Frequency. (b) Frequency. (c) Frequency. (d) Default. (e) Default. (f) Default.

Table 3 panel C and Fig. 3 panels (a) and (d) effectively implement this research design. There is a jump in the number of portfolio loans at 620 of 0.342 log points, and a jump in the default rate of these loans of 6.5 percentage points. Using Keys et al. (2010b)'s "unconditional probability" approach, one would interpret the jump in number of portfolio loans as a jump in

Table 4
Discontinuities in subgroups: jumbo, conforming, and 2008–2009 loans.

	FICO 620			FICO 660		
	log (*# loans*) (1)	*Delinquent* (2)	*Securitization* (3)	log (*# loans*) (4)	*Delinquent* (5)	*Securitization* (6)
Panel A: Jumbo Loans 2003–2007 (N=1,622,208)						
Change at cutoff	0.609***	0.034***	0.020**	0.307***	0.033***	0.006
s.e.	(0.012)	(0.010)	(0.008)	(0.008)	(0.005)	(0.005)
Base rate	–	0.312	0.807	–	0.247	0.791
Panel B: Conforming Loans 2003–2007 (N=12,875,935)						
Change at cutoff	0.307***	0.033***	−0.006***	0.102***	0.021***	−0.002***
s.e.	(0.003)	(0.002)	(0.001)	(0.003)	(0.001)	(0.001)
Base rate	–	0.232	0.935	–	0.153	0.934
Panel C: All Loans 2008–2009 (N=3,437,910)						
Change at cutoff	0.372***	0.007*	.002	0.210***	0.018***	0.007***
s.e.	(0.006)	(0.004)	(0.002)	(0.005)	(0.002)	(0.001)
Base rate	–	0.134	0.974	–	0.071	0.966

Notes: *Delinquent* is an indicator for whether the loan ever became 61 days delinquent within 36 months of origination for the 2003–2007 samples, and 18 months for the 2008–2009 samples. Securitization is an indicator for whether the loan was ever sold within 36 months of origination for the 2003–2007 samples, and 18 months for the 2008–2009 samples. The sample in Panel A includes only loans originated in 2003–2007 for amounts of above the GSEs' conforming loan limits. The sample in Panel B includes only loans originated in 2003–2007 for amounts below the GSEs' conforming loan limits. The sample in Panel C includes only loans originated in 2008–2009. Columns 1 and 4 use a local linear regression to test for a discontinuity in the number of loans at the credit score cutoff, as outlined in McCrary (2008). Columns 2, 3, 5, and 6 fit a 6th-order polynomial in FICO on either side of the cutoff. "Base rate" refers to the level of the dependent variable right below the credit score cutoff. Heteroskedasticity-robust standard errors in parentheses.

* Significant at 10%.

** Significant at 5%.

*** Significant at 1%.

originators' incentives to screen loans. Hence, the corresponding jump in defaults would be interpreted as evidence for the hypothesis that a *lack* of securitization led to lax underwriting. The problem with this interpretation of the jump in portfolio loans at 620 is that lenders changed their overall lending rate at this cutoff. This can be seen in the jumps in the number of private-label securitized loans and GSE securitized loans from the same origination years, also displayed in Fig. 3.

4.3.2. Isolating a private-label market: jumbo loans

Next we investigate whether the pattern of discontinuities in lender screening in the absence of discontinuities in securitization exists for a sample of loans at risk only of private-label securitization (and not GSE securitization). The GSEs are legally prohibited from buying loans for amounts larger than a certain size, called the conforming loan limit.[16] Loans larger than this limit are referred to as "jumbo." Because GSEs are prohibited from buying them, jumbo loans are only securitized by private-label securitizers. A straightforward way to isolate a purely private-label market is to select only jumbo loans.

Table 4 presents estimates of lending, default, and securitization discontinuities for jumbo loans in Panel A. Panel B presents similar estimates for conforming loans, the reciprocal category, which is dominated by GSE securitization. There are large jumps in the number of jumbo loans as well as their default rate. The number of loans roughly doubles at 620—there are 1530 loans in the sample with a FICO score of 619, compared with 3238 loans in the sample with a FICO score of 620—and the jump in number of loans at 660 is also large. The estimated default discontinuities are 3.4 and 3.2 percentage points at 620 and 660, respectively. There is a small jump in the securitization rate in the jumbo sample at 620, estimated at 2.0 percentage points from a base rate of 80.7% below 620. However, examination of the scatter plot in Fig. 4(c) reveals no visible securitization discontinuity at 620. At 660 there is no visible, nor statistically significant, discontinuity in the securitization rate.

A jump in securitization at 620 is potentially consistent with both the securitization rule-of-thumb theory and the origination rule-of-thumb theory. However, the increase in the securitization rate from 80.7% to 82.7% at 620 in the jumbo market seems too small to plausibly account for the large change in lender screening reflected by the *doubling* in the number of loans at 620. Thus even the 620 results are suggestive of a failure of the exclusion restriction. Moreover, the small change in the securitization rate at 620 could be the effect, not the cause, of changes in lender screening. The more intensive screening by lenders below 620 results in lenders having greater private information about those loans, as well as other changes in the characteristics of loans originated, and these differences may result in a lower fraction of loans being sold.

[16] For 2006 and 2007 in the continental U.S. the limit was $417,000.

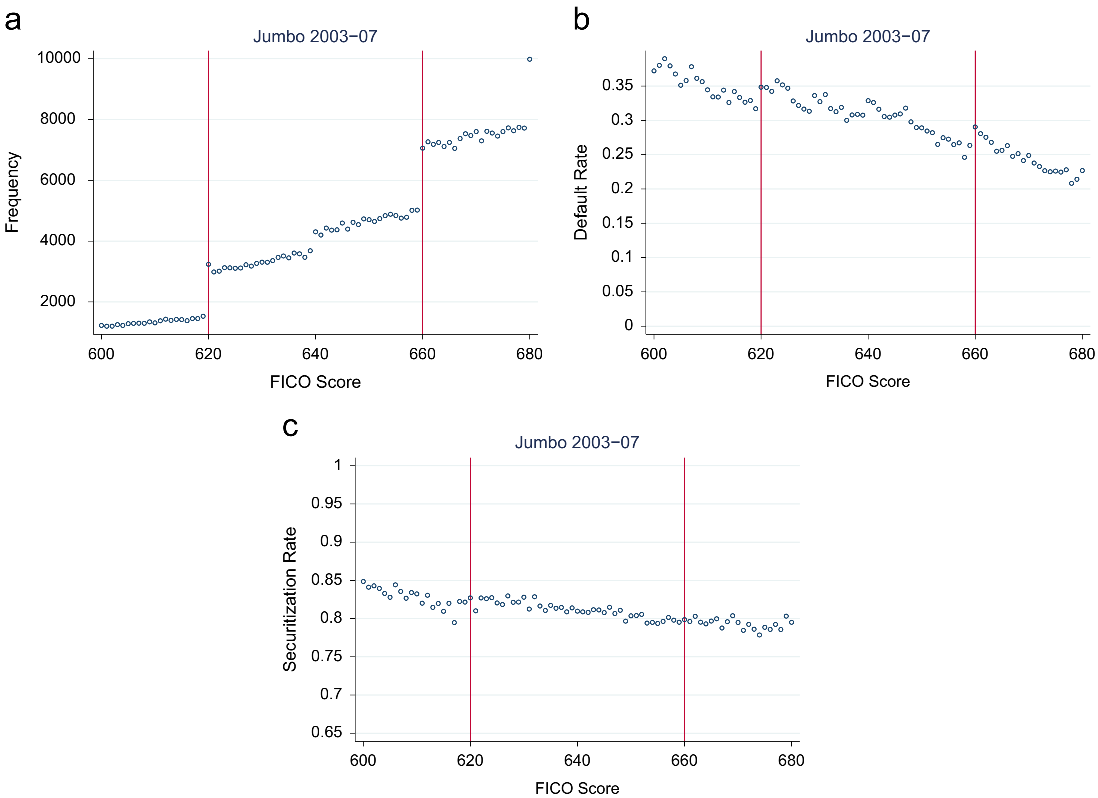

Fig. 4. Jumbo sample. Data source: Lender Processing Services Applied Analytics, Inc. Sample of first-lien, non-buydown, owner-occupied, single-family, jumbo mortgage loans originated between January 2003 and December 2007. A loan is considered defaulted if it ever became 61 days delinquent within 36 months of origination. A loan is considered securitized if it was ever sold within 36 months of origination: (a) Frequency. (b) Default. (c) Securitization.

4.3.3. Isolating a GSE market: loans originated in 2008-2009

One can isolate a sample of loans only at risk of GSE securitization by examining loans originated in 2008–2009, after the private-label MBS market shut down. During the peak years of 2005 and 2006, private-label securitizers issued $1191 billion and $1146 billion of MBS, respectively.[17] However, after the mass downgrades of private-label MBS by the rating agencies in 2007 that followed the onset of the subprime mortgage crisis, private-label MBS issuance ground to a halt. In 2007 only $707 billion private-label MBS were issued, followed by a mere $58 and $60 billion in 2008 and 2009, respectively. Of these, only $2 billion and $0.9 billion, respectively, were subprime MBS. This is negligible relative to the $1500 billion and $1800 billion in mortgages originated in 2008 and 2009, respectively. If credit score cutoff rules in screening are a product of a rule of thumb among private-label securitizers, as argued by Keys et al. (2010a, 2012), then there should be no discontinuities in our sample of loans originated in 2008–2009.

The 2008–2009 sample had a very different price appreciation experience than the 2003–2007 sample, since its loans were originated after the bursting of the housing bubble. Moreover, there is a much shorter follow-up period of only 18 months available in the data. These differences should result in smaller measured jumps in default for the same jump in borrower creditworthiness than in the 2003–2007 sample. Our ability to observe jumps in lending volume is unchanged.

Panel C of Table 4 and Fig. 5 reveals that the cutoff rules in screening at 620 and 660 FICO persist in this period without private-label securitization. There are large jumps in the number of loans at 620 and 660, showing that lenders continued to observe these cutoffs in their underwriting decisions. As expected, the jumps in default are relatively small but still statistically significant. Securitization appears smooth through the cutoffs, with no statistically significant jump at 620 and a jump of only 0.7 percentage points at 660. These results further confirm the origination rule-of-thumb theory and are inconsistent with the securitization rule-of-thumb theory.

4.3.4. Implications

Our institutional and quantitative evidence documenting the failure of the exclusion restriction implies that identifying a subset of loans for which there *is* a discontinuity in probability of securitization would not make the RD design valid. For example, Jiang et al. (2010) use data from a single originator and find a 2.7 percentage point jump in the securitization rate

[17] These statistics are from the 2011 *Mortgage Market Statistical Annual*, published by Inside Mortgage Finance.

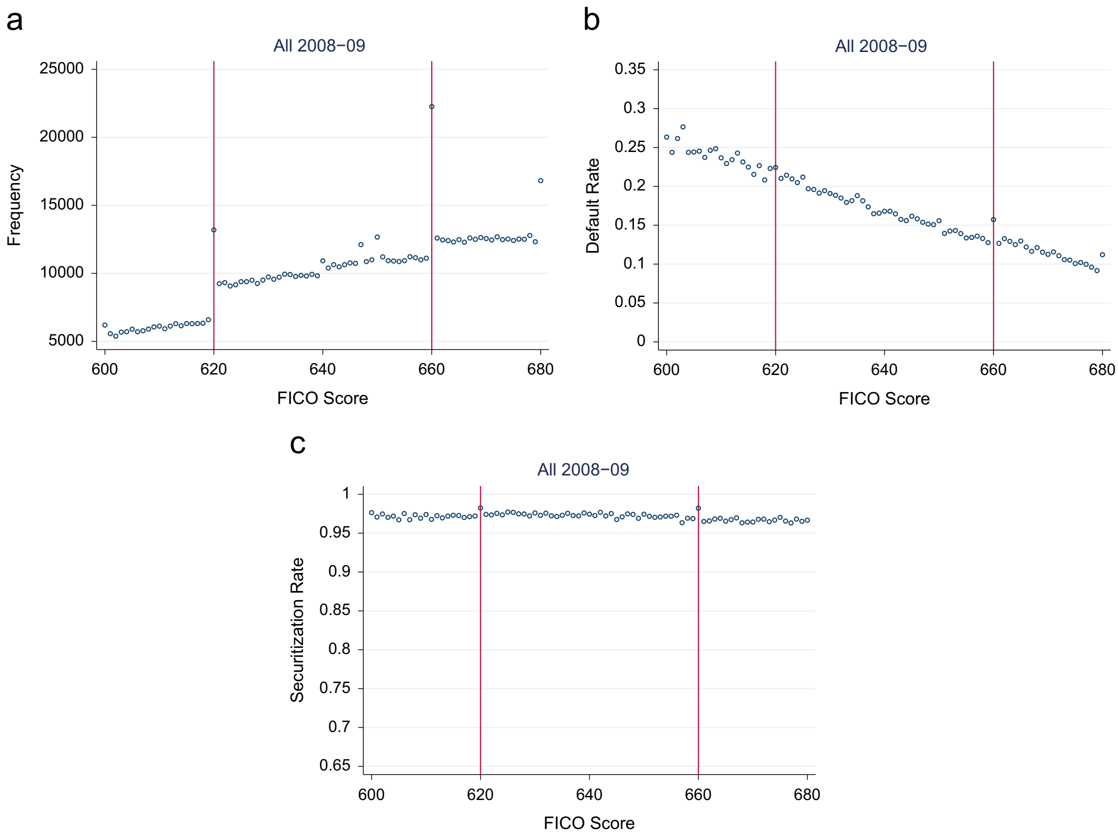

Fig. 5. 2008–2009 sample. Data source: Lender Processing Services Applied Analytics, Inc. Sample of first-lien, non-buydown, owner-occupied, single-family mortgage loans originated between January 2008 and June 2009. A loan is considered defaulted if it ever became 61 days delinquent within 18 months of origination. A loan is considered securitized if it was ever sold within 18 months of origination: (a) Frequency. (b) Default. (c) Securitization.

at 620, from a base of over 80%. They also find a 5.6 percentage point jump in default at 620, from a base of about 43%, which they interpret as the causal effect of securitization. Implicit in this research design is the assumption that if there were no moral hazard problem then there would not be a jump in default in their sample. However, default jumps at 620 even for samples of loans where there is *no* jump in securitization. This failure of the exclusion restriction means that regression discontinuity cannot provide evidence on the causal hypothesis of interest.

4.4. Other analyses

Keys et al. (2010a, 2012) raise several other issues with our empirical analysis. We investigated each of the issues they raise, and none change our basic analytical approach or conclusions. A detailed analysis of those issues is provided in the online appendix. Our analysis of three of those issues is provided below: using state anti-predatory lending laws to test the origination rule-of-thumb theory, analyzing a subsample of only low-documentation loans, and using propensity score weighting to isolate loans only at risk of private securitization.

4.4.1. Using variation from anti-predatory lending laws

Keys et al. (2010b, pp. 341–344) explicitly considers our central hypothesis—that the 620 FICO score threshold was used by lenders for reasons unrelated to securitization—and attempts to test it by using variation induced by the passage of state anti-predatory lending laws in Georgia and New Jersey in 2002 and 2003, respectively. The Georgia Fair Lending Act (GFLA)[18] and the New Jersey Home Ownership Security Act (NJHOSA)[19] restricted a range of lending practices by mortgage originators and also made assignees of mortgages, such as securitizers, liable if the originator violated the laws' proscriptions. The authors argue that these laws provide a test of the origination rule-of-thumb theory, arguing that if originators use 620 as a “cutoff for screening unrelated to securitization, we expect the passage of these laws to have had no effect on the differential screening standards around the threshold” (p. 342). They argue that, in contrast, if the

[18] O.C.G.A. Section 7-6A-1, *et seq.*
[19] N.J.S.A. 46:10B-22, *et seq.*

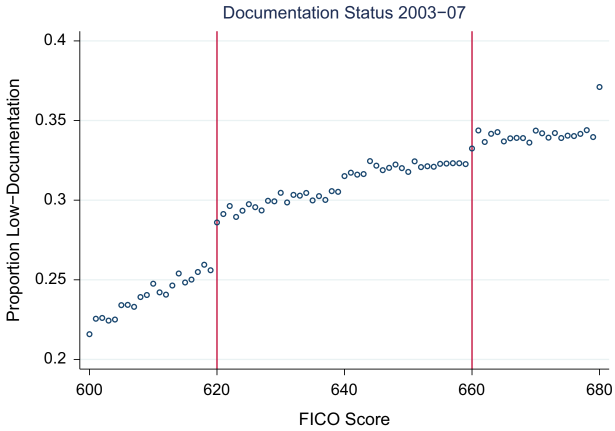

Fig. 6. Proportion low-documentation by FICO. Data source: Lender Processing Services Applied Analytics, Inc. Sample of first-lien, non-buydown, owner-occupied, single-family mortgage loans originated between January 2003 and December 2007. *Low-doc* includes both "low" and "no" documentation loans.

securitization rule-of-thumb theory is true and "these laws affect the differential ease of securitization around the [620] threshold … the screening differentials [around 620] we observed earlier should attenuate" (p. 342).

However, a key assumption of this test of the origination rule-of-thumb theory is that the laws changed only the ease of securitization, and that therefore the laws would have no effect on the size of the discontinuity in default at 620 under the origination rule-of-thumb theory. This assumption is inconsistent with the content of the laws. The goal of the GFLA was to prevent abusive lending practices. Georgia Senator Vincent Fort announced on January 24, 2001, that he would introduce the Georgia Fair Lending Act bill for the purpose of stopping "such lending practices as equity stripping, flipping, and packing, which are traditionally imposed by sub-prime lenders." Fort went on to discuss the risks associated with "predatory lending schemes" that most often target the "elderly, disabled veterans, 20 to 30-year-olds who are just starting out, and others who are financially strapped."[20] The purpose of the NJHOSA, similar in content to the GFLA, was also to prevent abusive lending practices. Accordingly, these laws regulated mortgage lending practices far more broadly than merely altering the ease of securitization, subjecting mortgage lending to a broad range of limitations and prohibitions.[21]

These laws can thus be expected to change the lending rate and default rate discontinuities at 620 under the origination rule-of-thumb theory and hence provide no way to test the two competing theories. The laws were designed to lower default levels directly. And because there are more high default risk loans above 620, the laws can be expected to have a bigger effect on default above 620 under both credit score cutoff theories. In contrast, our finding of jumps in lender screening in the absence of jumps in securitization rates unambiguously rejects the securitization rule-of-thumb theory.

4.4.2. Low documentation vs. full documentation

Keys et al. (2010a, 2012) also argue that we are inappropriately pooling together low-documentation and high-documentation loans in our analysis. Low-documentation (low-doc) loans, unlike standard full-documentation (full-doc) loans, require limited or no documentation of borrowers' income and assets. They argue that one should restrict analysis to low-doc loans because originators have more private information about low-doc loans than about high-doc loans, and hence the moral hazard problem is more severe.

However, a fundamental problem with selecting a sample according to documentation status is that it involves selecting on the outcome of interest, which is econometrically incorrect. The amount of documentation required by lenders is an important aspect of lender screening, which is the outcome variable of interest. Indeed, documentation status is a more direct measure of lender screening intensity than is loan performance. Fig. 6 plots the fraction of loans in our main sample that are low-doc loans as a function of FICO score. Consistent with other evidence that lender screening relaxes at cutoffs

[20] Tom Crawford, "Fort Introduces 'Georgia Fair Lending Act' Bill," *Georgia Report*, January 24, 2001, available at http://gareport.com/story/2001/01/24/fort-introduces-georgia-fair-lending-act-bill/.

[21] For example, the statutes prohibited creditors from making a home loan that finances various forms of insurance such as life, accident, health, or loss-of-income insurance. O.C.G.A. Section 7-6A-3(1)(B); N.J.S.A. 46:10B-25(4)(a). Both statutes prohibit creditors from knowingly or intentionally engaging in the "unfair act or practice of 'flipping' a home loan," defined as the extension of a home loan by a creditor to a borrower that refinances an existing home loan consummated within the prior 5 years when the loan "does not provide reasonable, tangible net benefit to the borrower considering all of the circumstances." O.C.G.A. Section 7-6A-4(a); N.J.S.A. 46:10B-25(4)(b). Both statutes prohibited creditors or servicers from recommending or encouraging default on an existing loan or other debt. O.C.G.A. Section 7-6A-3(2); N.J.S.A. 46:10B-25(4)(c). These statutes prohibit late payment fees in excess of 5% of the amount of the payment past due and restricts the practice of charging multiple late fees for the same past due payment. O.C.G.A. Section 7-6A-3(3); N.J.S.A. 46:10B-25(4)(d)(1); N.J.S.A. 46:10B-25(4)(d)(3). Furthermore, for "high-cost" loans, these statutes imposed numerous restrictions including limits on the rate at which scheduled payments could increase on adjustable rate mortgages, negative amortization, interest rate increases upon default, and the financing of points and fees. O.C.G.A. Section 7-6A-5(1)-(15); N.J.S.A. 46:10B-25(a)-(l). Moreover, the laws prohibit lenders from making certain high-cost loans unless the borrower has received counseling on the advisability of the loan and the lender reasonably believed that the borrower would be able to make the payments on the loan. O.C.G.A. Section 7-6A-5 (6)-(7); N.J.S.A. 46:10B-26(5)(g).

such as 620 and 660, there are large jumps in the fraction of low-doc loans at 620 and 660. Accordingly, selecting the sample based on the outcome produces biased estimates—a problem referred to by statisticians as posttreatment selection bias (Frangakis and Rubin, 2002). Valid causal inference in a regression discontinuity design requires pooling together units with all values of the outcome.

In addition to this general econometric problem caused by selecting on the outcome, the selection of low-doc loans raises a specific selection concern: the change in lender screening at 620 may shift the relative costs and benefits to loan applicants of providing full documentation instead of low documentation. To give one example of this form of selection, suppose that for a certain type of riskier-than-average loan applicant, providing full documentation of income and assets is especially costly. If these applicants are above the 620 cutoff, where lenders are less cautious, they have a good chance of receiving a loan even if they do not provide full documentation. In contrast, below the 620 cutoff, where lenders are more careful, it may become worth it to these applicants to fully document their income and assets for the lender. When considering a sample of only low-doc loans, this shift in the composition of which applicants provide low documentation may confound changes in the default rate stemming from lender screening itself.

Finally, as a theoretical matter, one would expect originators to have more private information about high-doc loans than about low-doc loans, the opposite of what Keys et al. suppose. For high-doc loans, lenders do a relatively thorough investigation of the borrower's income and assets, for example by calling the borrower's employer, investigating whether part of the downpayment is a gift versus a loan, and the like. In the process the lender accumulates private information about the borrower. In contrast, for low-doc loans the borrower simply states his income and assets with much less verification by the lender. Conceptually low-doc loans are loans for which the *borrower* has more private information vis-a-vis the lender. The asymmetric information problem that securitization raises, in contrast, has to do with private information that the *originator* has vis-a-vis the securitizer and investors. Thus it is not clear theoretically why one should expect the moral hazard problem to be more severe for the low-doc subgroup.

Nonetheless, analysis of low-doc and full-doc subgroups is included in the online appendix. There are jumps in lender screening in the absence of jumps in securitization for both low-doc and full-doc subgroups.

4.4.3. Propensity score weighting cannot isolate loans only at risk of private securitization

Keys et al. (2010a, 2012) use an alternative approach to our approach of examining the jumbo market in an attempt to isolate loans that were only at risk of private securitization and not GSE securitization. The authors assert that there are two separate markets: (1) the "non-agency" market of loans only at risk of being purchased by private securitizers and not by the GSEs; and (2) the "agency" market of loans only at risk of being purchased by the GSEs and not by private-label securitizers. They assume that loans are either in one market or the other and that no loan is both at risk of being bought by private-label securitizers and by the GSEs. Based on these premises, the authors separate loans into subgroups representing the agency and non-agency markets. The authors assign loans bought by private-label securitizers to the non-agency market and assign loans bought by the GSEs to the agency market. They use a complicated propensity score weighting approach to assign portfolio loans to one market or the other.

However, as discussed more formally in Appendix B in the online appendix, key to this approach is Keys et al.'s assumption that the GSEs and private securitizers operated in completely separate markets. This assumption is inconsistent with the evidence.

The LPS data provide a straightforward way to test this separate markets assumption. In particular, the data provide the entire history of ownership of each loan in the sample. It turns out that many loans that are ultimately owned by the GSEs are owned earlier by private-label securitizers. Conversely, many loans that were ultimately owned by private-label securitizers were earlier owned by the GSEs. Specifically, 18% of conforming loans in our 2003–2007 sample were at one point in their life held by the GSEs and at another point owned by private-label securitizers. This is a lower bound for the fraction of loans in the sample that were at risk of being sold to both the GSEs or private-label securitizers, since loans that were *ex ante* at risk of being sold to both types of buyers need not have been held by both types *ex post*.

Moreover, institutional evidence also reveals that the GSEs actively competed with private-label securitizers to buy subprime loans. Calomiris and Wallison (2008) write (emphasis added):

> In order to curry congressional support after their accounting scandals in 2003 and 2004, Fannie Mae and Freddie Mac committed to increased financing of "affordable housing." They became *the largest buyers of subprime and Alt-A mortgages* between 2004 and 2007, with total GSE exposure eventually exceeding $1 trillion.

As this quotation makes clear, the GSEs were large buyers of subprime mortgages during the period in question. This is also reflected in reports from the GSEs' federal regulator. For example, Office of Federal Housing Enterprise Oversight (2007, p. 11) reports that, in addition to being major buyers of subprime private MBS, Fannie Mae and Freddie Mac "also buy and hold subprime mortgages directly." Moreover, the SEC filings of subprime originators show that they sell subprime loans to both the GSEs and private-label securitizers. For example, Countrywide Financial, one of the largest originators of subprime mortgage loans, reported in its 2005 Annual Report that, in addition to pooling "Nonprime Mortgage Loans" into private-label MBS, it "pooled a portion of [its] Nonprime Mortgage Loans into securities guaranteed by Fannie Mae" (p. 95).

Finally, as discussed above, the failure of the exclusion restriction means that even if one can identify a subset of loans that were only at risk of private-label securitization for which there is a jump in the securitization rate at 620, the RD research design would not be valid for that subset. It would be impossible to know what part of any jump in default is due to the moral hazard effect of securitization, and what part is due to the failure of the exclusion restriction.

4.4.4. Lenders that sell all of their mortgages disprove the securitization rule-of-thumb theory

Finally, Keys et al. (2010b) perform a robustness check by selecting loans from independent mortgage companies that do not keep a portfolio of loans on their books. For this subgroup of loans, they find similar jumps in default at 620. They interpret this robustness check as showing that their results are not being driven by selection of loans into securitization versus being held in portfolio. But theoretically, the moral hazard hypothesis being tested under the securitization rule-of-thumb theory is that lenders screen less carefully loans with a higher probability of being sold. For this subgroup of loans, every loan, both above 620 and below, has the same probability of securitization, namely 1. Hence the change in lender screening at 620 demonstrated by Keys et al. (2010b) is inconsistent with the securitization rule-of-thumb theory and consistent with the origination rule-of-thumb theory.

5. Conclusion

This paper shows that lenders adopted credit score cutoff rules directly in response to underwriting guidelines from Fannie Mae and Freddie Mac. Institutional evidence shows that the GSEs required lenders to adopt credit score cutoff rules and that these rules spread through their incorporation into underwriting software, eventually becoming de facto industry-wide origination standards. The GSEs established these credit score cutoff rules in order to improve underwriting. Evidence from a loan-level dataset shows that there are lender screening cutoffs at 620 and 660 in the absence of a discontinuous increase in securitization. The same pattern exists among jumbo loans, a purely private-label market. The data show that these lender screening cutoffs persisted in 2008–2009, after the private-label securitization market had shut down.

This institutional and quantitative evidence shows that credit score cutoff rules unfortunately do not provide a useful laboratory for estimating the moral hazard effect of securitization on lender screening. Lenders change their screening behavior at credit score cutoffs for reasons other than a change in the probability of securitization. Therefore the change in default at these credit score cutoffs is not evidence that securitization led to lax screening by lenders.

The cutoff rule evidence does tell us, however, that Fannie Mae and Freddie Mac were to a significant extent successful in implementing their desired underwriting guidelines throughout the mortgage industry. After concluding that credit score cutoff rules in screening would improve mortgage underwriting, Fannie and Freddie used contractual provisions, monitoring, and software systems to ensure that originators adopted them. Instead of showing that securitization led to moral hazard, the cutoff rule evidence thus provides some evidence for the opposing hypothesis: securitizers were to a significant extent able to control lenders' underwriting behavior.

There are also additional conceptual problems with the theory that a moral hazard problem caused by securitization was a significant contributor to the recent financial crisis. Information economics provides a well-developed theory for the incentive problems posed by information asymmetry. The main implication of this theory is that, in equilibrium, incentive problems inhibit trade. If moral hazard were a major problem in the securitization market, then, the standard incentive theory would predict that relatively few mortgages would be securitized. Investors would require a higher interest rate to compensate them for the increased credit risk of mortgages caused by securitization, and this would result in less trade in mortgages. In fact, the vast majority of mortgages were securitized in the run-up to the financial crisis. So the basic stylized facts of the crisis are inconsistent with the standard moral hazard theory.

One possible way to resurrect the moral hazard explanation of the mortgage crisis is to depart from the neoclassical model and suppose that investors were naive about the incentive problems posed by securitization. Under this account, investors would not require compensation for the higher credit risk of mortgages caused by securitization, and hence in equilibrium the moral hazard problem would not inhibit trade but would increase the riskiness of mortgages. However, Gerardi et al. (2008) and Foote et al. (2012) document that investors formed accurate views of the risks posed by the decline in underwriting standards in the run-up to the crisis and made accurate predictions of the performance of mortgage-related securities conditional on particular assumptions about housing price appreciation.

Investors' big mistake, then, was not in misunderstanding a moral hazard problem or underwriting problems in general. Rather, it was in putting too little weight on the possibility that housing prices would fall—a classic bubble phenomenon. For example, Foote et al. (2012) describe a report by Lehman Brothers from August 2005 that put only a 5% probability of housing prices declining nationally.

Unfortunately, the reforms to mortgage securitization adopted since the crisis have been premised on the view that moral hazard due to securitization, rather than the housing bubble, was the main cause of the mortgage crisis, and to a large extent this view is based on findings that rely on the securitization rule-of-thumb theory to interpret the evidence. We hope that our analysis showing that this theory is incorrect will lead policymakers to reevaluate their approach.

Acknowledgments

Financial support for this research was provided by the John M. Olin Center for Law, Economics, and Business at Harvard Law School. We thank Viral Acharya, Yakov Amihud, Jennifer Arlen, Ken Ayotte, Michal Barzuza, Effi Benmelech, Stephen

Choi, Andrew Eggers, Chris Foote, Josh Gallin, Ed Glaeser, Claudia Goldin, Jeff Gordon, Robin Greenwood, Andrew Haughwout, Kate Judge, Marcel Kahan, Larry Katz, Benjamin Keys, David Laibson, Chris Mayer, Justin McCrary, Doug McManus, David Scharfstein, Josh Schwartzstein, Amit Seru, Andrei Shleifer, Paul Smith, Holger Spamann, Jeremy Stein, Adi Sunderam, Joseph Tracy, James Vickery, Vikrant Vig, Glen Weyl, Larry White, Paul Willen, Heidi Williams, Noam Yuchtman and numerous seminar participants for valuable comments and discussions. We thank Larry Cordell, Vidya Shenoy, and Mark Watson for their patient help with the LPS dataset. We are grateful to the Research Departments at the Federal Reserve Bank of Boston and the Federal Reserve Bank of New York for supporting this research. We thank Xiaoqi Zhu for outstanding research assistance. The analysis and conclusions are the authors' own and do not indicate concurrence by other members of the Federal Reserve System research staff or the Board of Governors.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at http://dx.doi.org/10.1016/j.jmoneco.2014.01.005.

References

Akerlof, G.A., 1970. The market for lemons: quality uncertainty and the market mechanism. Q. J. Econ. 84 (3), 488–500.
Calomiris, C.W., Wallison, P.J., 2008. Blame Fannie Mae and Congress for the Credit Mess. Wall Street J, Sept. 23, 2008. (http://online.wsj.com/news/articles/SB122212948811465427).
Ellison, G., Holden, R., 2008. A Theory of Rule Development. Unpublished manuscript.
Fannie Mae, 1995. LL09-95: Measuring Credit Risk: Borrower Credit Scores and Lender Profiles. Letter to Lenders.
Fannie Mae, 1997. LL01-97: Mortgage Underwriting Tools–Automated Underwriting and Credit Scores: Measuring Credit Risk: Borrower Credit Scores and Lender Profiles. Letter to Lenders.
Fannie Mae, 2007. Guide to Underwriting with DU. Letter to Lenders.
Foote, C.L., Gerardi, K., Willen, P.S., 2012. Why Did so Many People Make so Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis. NBER Working Paper 18082.
Frangakis, C.E., Rubin, D.B., 2002. Principal stratification in causal inference. Biometrics 58 (1), 21–29.
Freddie Mac, 1995. The Predictive Power of Selected Credit Scores. Industry Letter.
Gerardi, Kristopher S., Andreas Lehnert, Shane M. Sherlund, Willen, Paul S., 2008. Making sense of the subprime crisis. Brookings Pap. Econ. Act. 2, 69–145.
Gorton, G.B., 2009. The subprime panic. Eur. Financ. Manag. 15 (1), 10–46.
Hutto, G., Lederman, J., 2003. Handbook of Mortgage Lending. Mortgage Bankers Association of America.
Jiang, W., Nelson, A.A., Vytlacil, E., 2010. Securitization and Loan Performance: A Contrast of Ex Ante and Ex Post Relations in the Mortgage Market. Unpublished Manuscript.
Keys, B., Mukherjee, T., Seru, A., Vig, V., 2010a. 620 FICO, Take II: Securitization and Screening in the Subprime Mortgage Market. Unpublished Manuscript.
Keys, B.J., Mukherjee, T., Seru, A., Vig, V., 2009. Financial regulation and securitization: evidence from subprime loans. J. Monet. Econ. 56 (5), 700–720.
Keys, B.J., Mukherjee, T.K., Seru, A., Vig, V., 2010b. Did securitization lead to lax screening? Evidence from subprime loans. Q. J. Econ. 125 (1).
Keys, B.J., Seru, A., Vig, V., 2012. Lender screening and the role of securitization: evidence from prime and subprime mortgage markets. Rev. Financ. Stud. 25 (7), 2071–2108.
Krainer, J., Laderman, E., 2009. Mortgage Loan Securitization and Relative Loan Performance. Unpublished manuscript.
McCrary, J., 2008. Manipulation of the running variable in the regression discontinuity design: a density test. J. Econ. 142 (2), 698–714.
Office of Federal Housing Enterprise Oversight, 2007. 2007 Performance and Accountability Report.
Quinn, L., 2000. Credit score scrutiny. Mortg. Bank. 60 (12), 50–55.
Rajan, U., Seru, A., Vig, V., 2010. The Failure of Models that Predict Failure: Distance. Incentives and Defaults. Unpublished Manuscript.