A Bayesian analysis of payday loans and their regulation[☆]

Mingliang Li [a,*], Kevin J. Mumford [b], Justin L. Tobias [c]

[a] SUNY-Buffalo, Department of Economics, 439 Fronczak Hall, Buffalo, NY 14260, United States
[b] Purdue University, Department of Economics, 100 S Grant Street, West Lafayette, IN 47907, United States
[c] Purdue University, Department of Economics, 403 W State Street, West Lafayette, IN 47907, United States

ARTICLE INFO

ABSTRACT

Payday loans are small short-term loans that a borrower must repay or renew on his/her next payday. In states where payday lending is legal, many terms of these loans are regulated, ostensibly to protect the consumer from excessively burdensome lending practices.

The existing literature on payday loans has primarily focused on estimating causal effects of access to those loans, including work by Morse (2011), Skiba and Tobacman (2009) and Melzer (2011). Using individual-level administrative records on borrowers in 38 states from an online payday lender, this paper departs from past work by estimating how payday loan regulation affects borrower behavior, specifically how much they choose to borrow, how many times they choose to renew the loan, and whether or not they choose to default. State-level variation in maximum loan sizes and renewal caps are used as exclusion restrictions for identification purposes. We pay particular attention to the calculation of posterior predictive distributions that summarize the sensitivities of borrower behavior to various changes in state-level policies.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Payday loans are short-term loans made to individuals that are either repaid or renewed on the borrower's next payday. These loans typically carry a 14-day loan term and average approximately $300. Although the amount of interest paid on such a loan is often reasonably small (given the small loan principals), interest rates charged on such loans are typically very high, with Annual Percentage Rates (APRs) often in excess of 400%. Given the magnitudes of these rates, the practice of payday lending is regarded as predatory by many, and the industry has come under increased scrutiny of late, with some states effectively outlawing payday lending through existing usury legislation.[1]

In states where online payday lending is legal, many terms of the loan are regulated. States impose limits on the interest rate charged, the maximum amount that can be borrowed, the length of the loan and what penalties can be imposed on borrowers who default. Averaged across different states, a typical interest rate limit is approximately 15% for a two-week loan, meaning that the

borrower must pay $15 in interest on a $100 loan at the end of the two week period. States regulate if a borrower is allowed to renew at the next payday and extend the loan period and, if so, how many times a borrower is allowed to renew. Borrowers that renew their loan pay the interest accumulated to that point and may also choose to pay down some of the principal.

How these state policies affect borrowers' behavior is not well understood in the literature and in this paper we aim to better understand these relationships. For example, if a state were to reduce the maximum interest rate that can be charged, will borrowers respond by paying back their loans more quickly or will the lower interest rate cause them to delay repayment? Is it possible that payday lenders may not strongly object to the setting of lower interest rates if it turns out that lower rates are associated with lower probabilities of default, increased loan amounts and longer loan durations? In this paper we seek to address these and other questions by modeling the borrowers' choices of loan amount, loan length, and whether or not to default, and determining how these depend on state-level policies. The data we employ in our analysis come from a single online payday lender with loans in 38 states over a fairly short period of time.

The economic literature on payday lending has primarily focused on estimating causal effects of *access* to payday loans on various outcomes, including work by Morse (2011), Skiba and Tobacman (2009), and Melzer (2011). The evidence produced by these studies is mixed, with Morse (2011) finding that access to payday loans decreases foreclosures while Skiba and Tobacman

---

[1] In Georgia, for example, payday lending is effectively prohibited through existing legislation regarding permissible interest rates on small loans.

(2009) find that access to payday loans increases the likelihood of bankruptcy. In a similar spirit of questioning the value of payday lending, Melzer (2011) finds that access to payday loans decreases the borrower's likelihood of paying mortgage, rent, and utility bills. Other papers including Flannery and Samolyk (2005) and Skiba and Tobacman (2007) have been concerned with reconciling how payday lenders that charge interest rates of 15% for a two-week loan (391% APR) make only normal levels of profit with annual returns on equity around 10%. They find that the large number of borrowers who default on their loans is the major factor that drives down the average return on equity for payday lenders.

Although the references listed above examine potential benefits or pitfalls of the payday loan industry, few studies in this literature have specifically examined the impact of payday loan regulation on the behavior of borrowers. Skiba and Tobacman (2008), perhaps the closest study to the one performed here, estimate a structural model of repayment and default behavior in order to distinguish between alternative models of borrower preferences. They find that quasi-hyperbolic discounting models perform better than exponential discounting models. However, they neglect to perform model simulations under alternative policies and thus do not comment on how changes to payday loan regulation would likely influence the types of borrower behavior that they consider.

In this paper we use cross-state differences in payday loan regulation to identify the effects of these policies on borrower behavior. The state-level policies we consider are the maximum interest rate, the maximum loan amount, the maximum number of times a loan can be renewed, and the penalty that can be imposed on those who default. We estimate the effects of these policies on the borrower's decisions of how much to borrow, how long to keep the loan, and whether or not to default.

The model we employ to accomplish this goal is a three-equation nonlinear triangular simultaneous equation model. Our specification flexibly represents the latent outcomes as being generated by a finite mixture of Gaussian distributions and the model is fit using MCMC methods. We uncover a number of relationships, finding, for example, that interest rate reductions lead to longer loan durations and lower probabilities of default while reductions in the maximum amount that can be borrowed reduce the number of times loans are renewed and also lower the probability of default.

Although our paper is applied in nature, the general theme of Bayesian modeling of SEMs is consistent with several influential works of Herman van Dijk (e.g., Kloek and van Dijk, 1978, Zellner et al., 1988 and Kleibergen and van Dijk, 1998, among others), whom we honor here with this volume. Our work, though investigating a very different question, also shares the same general structure and strategy for posterior simulation as a number of other applied Bayesian systems of latent equations analyses, including Li (1998), Geweke et al. (2003), Munkin and Trivedi (2003, 2008) and Deb et al. (2006), to name a few.

The outline of this paper is as follows. We describe the data in Section 2, while the econometric model is introduced and discussed in Section 3. Model diagnostics and empirical results are provided in Section 4 and the paper concludes with a summary in Section 5.

2. Data

The data come from a single online payday loan lender who began making online loans in the fall of 2006. After approximately two and a half years of operation and issuing approximately 2500 loans, our lender promptly stopped online advertising. This substantially reduced the number of new loans and provided us with an opportunity to see the existing loans through to repayment

Table 1
State-level payday loan policies.

Panel A: maximum interest rate		
Interest rate	Number of loans	Percent of loans
0.1	199	8.57
0.15	773	33.28
0.155	79	3.40
0.1675	34	1.46
0.175	37	1.59
0.18	261	11.24
0.2	131	5.64
0.23	809	34.83

Panel B: maximum loan amount		
Max amount	Number of loans	Percent of loans
300	245	10.55
350	55	2.37
400	12	0.52
500	620	26.69
550	47	2.02
600	129	5.55
700	14	0.60
800	119	5.12
900	56	2.41
1000	1026	44.17

or default without new loans being added to the data. Of the 2500 loans about 700 ended in default.

Borrowers likely learned of our online lender by doing an online search of the term "payday loan" or "quick loan". They start the loan process by going to the lender's website and selecting their state of residence. The borrower is then directed to an online form to specify the desired amount of the loan, the date of the next payday, income, contact information, and employer contact information. The form states that the borrower must be employed with the same company for 3 consecutive months and must have take-home pay of at least $1000 per month (approximately full-time work at the minimum wage) to be approved. In practice, however, our lender generally made no effort to contact the employer or to verify the reported income.[2]

Borrowers came from 38 states with substantial heterogeneity in payday loan regulation. As shown in Table 1, the interest rate ranged from 10% to 23% depending on the state of residence. The lender applied a 23% interest rate (600% APR) in states that set a maximum rate higher than 23% and in states with no regulation limiting the interest rate. The maximum amount that can be borrowed ranged from $300 to $1000 (also shown in Table 1 with the lender's $1000 cap applying in states that have a higher limit or have no limit to the amount that can be borrowed).

When a payday loan is made, it is due on the next payday. About 80% of the borrowers in the data are paid twice a month or every two weeks and generally have a loan length of about 14 days. Borrowers with a longer period of time between paydays are given longer to repay the loan without having to pay additional interest charges. The interest rates reported in Table 1 are for the loan period, regardless of its length. Borrowers who pay back the full principal within 3 days (or 4 in some states) have their loan canceled and do not owe any interest. Once this grace period has been exceeded, borrowers do not receive a reduction in interest charged for early repayment.

The lender will allow the borrower to renew the loan up to five times if this is allowed by the state. To renew the loan, the

---

[2] The borrower's name, social security number, address, and checking account number were checked against a national database in order to reject applications from those who have a history of defaulting on payday loans. The lender accepted nearly every application that passed this check and did not run a separate credit check or use the self-reported income and employment information in order to differentiate applications.

borrower must pay at least the amount of interest due and can choose to pay more to reduce future interest costs. In 18 of the states where the lender made loans, state law does not allow borrowers to renew loans in this way. However, online lenders often provide the borrower with a new loan to cover paying off the principal and in this way effectively renew the loan. We see evidence of this practice in the data and categorize it as a renewal of the existing loan, but suspect that the additional complication caused by needing to take out a new loan to pay off the old loan may have an effect on borrower behavior. In 7 of the 20 states where payday loan renewals are allowed, the state's limit on the number of renewals is less than five.

If a borrower defaults on the loan, the payday lender makes efforts to contact the borrower to encourage loan repayment. If these efforts are unsuccessful, the lender sells the loan to a collection agency. Payday loan default is not a criminal offense (although collection agencies sometimes claim that it is), and most states have placed limits on the amount of collection and attorney fees that can be collected in addition to the amount owed on the loan. The additional collection fee may have an influence on the individual's decision to default on the loan.

Of those loans that end in default, 55% never make any payments or renew the loan by paying the interest; we refer to this as immediate default. The remaining 45% make some payments before defaulting. While surprising, this is confirmed by Skiba and Tobacman (2008) who find that defaulting borrowers have on average repaid an amount equal to 90% of the loan's principal at the time they default.

For each loan we observe the amount of the loan, the interest rate, the term length, the number of renewals, if the loan ended in default, and the state of residence which allows us to match each loan to the state's maximum loan amount, number of renewals allowed, and default penalties. Unfortunately, the lender did not ask borrowers for demographic information like gender and race and did not retain the self-reported age and income information once a loan was made. However, we were able to obtain a mailing address for each borrower and match this with available public records to obtain the estimated monthly rent associated with the address.[3] Rent is a good proxy for the financial resources of an individual as rent is a major component of consumption and is highly correlated with income. We suspect that it is likely a better measure of financial resources than the unverified income that applicants report to the lender in order to qualify for a loan. We also used the address to link each loan to the median household income, average education level, and fraction of the population that are Black and Hispanic in the borrower's zip code.

Summary statistics for the data used in the analysis are reported in Table 2. For loans where the mailing address was not for the borrower's home (usually a PO box or work address), we were unable to obtain the monthly rent and drop those loans for this analysis. This resulted in a loss of only about 5% of the observations. Another 2% of the observations were dropped from the analysis because of other missing or miscoded values, leaving us with more than 2300 observations.

3. The model

In our analysis of the payday loan data, we imagine that the agent makes the following sequence of choices: First, she decides

Table 2
Summary statistics.

Variable	Obs	Mean	Std. dev.	Min.	Max.
Interest rate	2323	0.1806	0.0420	0.10	0.23
Loan amount	2323	$302.9	$161.5	$39.4	$1000
Maximum amount	2323	$728.5	$269.0	$300	$1000
Default	2323	0.2755	0.4469	0	1
Immediate default	2323	0.1550	0.3620	0	1
Default penalty	2323	$93.0	$79.6	$15	$330
Rent	2323	$967.4	$362.1	$367	$2388
Loan term	2323	15.52	4.74	3	31
Loan renewals	1963	2.434	1.834	0	5

Notes: The loan renewals is reported for those who are not immediate defaulters.

on the loan amount.[4] Second, conditioned on a loan amount, the agent decides how many times to renew the loan (that is, to make an interest payment and extend the duration of the loan) as well as whether or not to default. As stated in the previous section, we observe both a sizeable fraction of "immediate" defaulters (who never make any repayment whatsoever and thus may take out the loan with every intention of defaulting) as well as those who make several interest payments, only to eventually default. We construct our model to capture the existence of both types of defaulters and later will perform diagnostic checking to verify that our specification is able to reproduce these distinct patterns of behavior that are present in the data.

Given this assumed decision-making structure, the model we employ is a nonlinear triangular simultaneous equation model (SEM).[5] To allow for flexibility in modeling the distribution of our outcomes, we employ a finite Gaussian mixture model. We represent this below in the traditional way by augmenting the parameter space to include a series of component labeling variables $c_i \in \{1, 2, \ldots, G\}$, for $i = 1, 2, \ldots, n$, that denote component membership (that is, if $c_i = g$ then individual i "belongs to" the gth component of the mixture).

Conditioned on the values of these component labels, we express our model as:

$$LogAmt_i^* = \beta_{0c_i} + \beta_1 Rate_i + \beta_2 Term_i + \beta_3 LogRent_i + \beta_4 LogStateMAX_i + \epsilon_i \tag{1}$$

$$LogAmt_i = \min\{LogAmt_i^*, \ LogStateMAX_i\} \tag{2}$$

$$Renewals_i^* = \gamma_{0c_i} + \gamma_1 LogAmt_i + \gamma_2 Rate_i + \gamma_3 Term_i + \gamma_4 LogRent_i + \gamma_5 StateCap_i + \gamma_6 NoStateCap_i + u_i \tag{3}$$

$$Renewals_i = \sum_{j=0}^{5} j \times I(\delta_{j+1} < Renewals_i^* \le \delta_{j+2}) \tag{4}$$

$$Default_i^* = \theta_{0c_i} + \theta_1 LogAmt_i + \theta_2 I(Renewals_i = 1) + \theta_3 I(Renewals_i = 2) + \theta_4 I(Renewals_i = 3) + \theta_5 I(Renewals_i = 4) + \theta_6 I(Renewals_i = 5) + \theta_7 Rate_i + \theta_8 Term_i + \theta_9 LogRent_i + \theta_{10} LogStatePenalty_i + \theta_{11} StatuteLimit_i + v_i \tag{5}$$

$$Default_i = I(Default_i^* > 0). \tag{6}$$

For the error terms, we specify

$$\begin{pmatrix} \epsilon_i \\ u_i \\ v_i \end{pmatrix} \middle| \cdot, c_i \overset{\text{ind}}{\sim} N \left[\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \Sigma_{c_i} = \begin{pmatrix} \sigma_{\epsilon\epsilon c_i} & \sigma_{\epsilon u c_i} & \sigma_{\epsilon v c_i} \\ \sigma_{u\epsilon c_i} & \sigma_{u u c_i} & \sigma_{u v c_i} \\ \sigma_{v\epsilon c_i} & \sigma_{v u c_i} & \sigma_{v v c_i} = 1 \end{pmatrix} \right] \quad (7)$$

and the conditional prior probability that i will be drawn from the gth component of the mixture is denoted as η_g:

$$\Pr(c_i = g | \boldsymbol{\eta}) = \eta_g, \quad \text{for } g = 1, 2, \dots, G \quad \text{and} \quad \sum_{g=1}^{G} \eta_g = 1. \quad (8)$$

In Eqs. (1)–(7), we allow intercepts and covariance matrices to differ across mixture components, but impose equality of the regression slopes across components of the mixture. This restriction aids in interpreting our results, mitigates concerns regarding identification (as discussed in greater depth below) and also introduces parsimony in the number of parameters that are employed. By allowing intercepts and covariance matrices to vary across the components, however, we introduce a great deal of flexibility to the analysis. In the following section we will examine how well this specification is able to capture key features of our data.

Eqs. (1)–(2) describe the amount of the loan taken by the individual (*LogAmt*), with the latent loan amount potentially depending on the charged interest rate (*Rate*), the length of time between borrower paydays (denoted *Term*, which is typically 14 since most individuals are paid biweekly), the monthly rent paid by the individual (*LogRent*) and the established state maximum amount (*LogStateMax*).

Although it may seem inappropriate at first, we model the interest rate as exogenous. In this analysis, the interest rate of the loan is not determined by any type of negotiation between lender and borrower, or assessment by the lender of the agent's (unobservable) creditworthiness. Instead, our lender simply sets the interest rate equal to the state-imposed maximum, which is assumed to be exogenous.[6] The observed amount borrowed in (2) is then set equal to the smaller of the desired latent amount and the maximum loan amount that is established by the state.

In (1) we additionally include *LogStateMax* as an explanatory variable, allowing for the possibility that state loan caps can influence outcomes beyond just support restrictions. When completing the online loan application, the agent is informed of the maximum amount that can be borrowed in their state of residence. The appearance of such information may affect the borrower's amount decision independently of the bounds on the loan amount, and we include the (log) state maximum as a covariate to allow for this possibility. Finally the variable *Term* is added to capture possible heterogeneity between biweekly and monthly paid individuals and possible differences in borrowing behaviors across these groups.

Eqs. (3)–(4) model the number of times that the agent renews the loan, with the latent renewals variable potentially depending on the loan amount, interest rate, loan term, monthly rent paid by the agent, and state policies regarding the number of permissible renewals (*StateCap*, *NoStateCap*). As briefly discussed in the previous section, the state-set renewal cap is not really binding in practice (and certainly does not bind within our data), as the agent can circumvent the state cap by simply taking out a new loan to effectively extend the payoff deadline. Since some states in our data do not have a cap on the maximum number of loan renewals, the variable *StateCap* is not well-defined in such

cases. Our practice for these states is to simply set the state cap variable equal to zero and to define and separately include a dummy variable for the lack of a state renewal cap (*NoStateCap*).

Renewals is modeled as an ordinal response, with values ranging from 0 to 5, where a value of 0, for example, indicates that a non-defaulting agent will pay off the loan plus interest charges at the next payday. Our lender limits renewals at five, and thus we truncate the value of *Renewals* at this upper limit. The linkage between latent renewals and the observed renewals outcome is described by (4), with $\boldsymbol{\delta}$ denoting a vector of estimable cutpoint parameters.

Our approach for identifying and normalizing these cutpoints departs slightly from what is most commonly done in the literature. We begin with the usual restrictions, namely: $\delta_1 = -\infty$, $\delta_2 = 0$ and $\delta_7 = \infty$. However, instead of estimating the remaining cutpoint parameters in the traditional way, we restrict the largest unknown cutpoint, δ_6, to equal unity and commensurate with this, do not restrict the variance of the error in (3) to equal unity. This equivalent specification can be derived by first taking the latent equation in (3), multiplying it by $1/\delta_6$, and then interpreting the variance of the transformed error as the squared reciprocal of the largest unknown cutpoint. This produces an observationally equivalent ordinal model, as noted by Nandram and Chen (1996), with the desirable properties that: (a) a diagonal restriction on $\boldsymbol{\Sigma}_g$ is no longer necessary, which can greatly facilitate computation, (b) one fewer cutpoint needs to be dealt with (as posterior simulation of the cutpoint parameters is decidedly less standard) and, in the special case of an ordinal variable with three possibilities, no unknown cutpoints are present. Finally, as noted in Nandram and Chen (1996), (c) the rescaling transformation appears to offer improved mixing performance relative to the traditional scale-restricted approach. We describe more details regarding posterior simulation of the cutpoint parameters in the Appendix.[7]

In the default equation we include the loan amount, separate indicator variables for the number of renewals, the interest rate, loan term, and two additional variables (*LogStatePenalty*, *StatuteLimit*) that may influence the decision to default. The *LogStatePenalty* variable is pieced together from state-level policies regarding what the lender can legally seek to collect from an agent who has defaulted on a $300 loan. These include fees in various proportions to the size of the loan in addition to "reasonable" attorney fees that, at least for some states, are refundable to the lender.[8] Finally, *StatuteLimit* is a state-level variable that indicates how long the lender (or a collection agency) legally has to collect on the debt. Once this limit is reached, debt collectors may no longer sue the borrower.

The model is completed by choosing priors of the forms

$$\boldsymbol{\psi} \sim \mathcal{N}(\boldsymbol{\psi}_0, \boldsymbol{V}_\psi) \quad (9)$$

$$p(\boldsymbol{\Sigma}_g) \propto p_{IW}(\boldsymbol{\Sigma}_g | p\boldsymbol{R}, p) I(\sigma_{vvg} = 1), \quad g = 1, 2, \dots, G \quad (10)$$

$$\boldsymbol{\eta} = (\eta_1, \eta_2, \dots, \eta_G) \sim \text{Dirichlet}(\kappa_1, \kappa_2, \dots, \kappa_G) \quad (11)$$

$$p(\boldsymbol{\delta}) \propto I(\delta_1 = -\infty) I(\delta_2 = 0) I(0 < \delta_3 < \delta_4 < \delta_5 < 1) \\ \times I(\delta_6 = 1) I(\delta_7 = \infty), \quad (12)$$

---

[6] Our lender, however, does impose a rate of 0.23 as a universal maximum: if the state-imposed maximum exceeds 23%, the lender still sets the interest rate equal to 0.23.

[7] Li and Tobias (2008) provide more details on this rescaling transformation, while Jeliazkov et al. (2008) consider it, along with other possibilities, in the context of multivariate ordinal responses.

[8] The collection agency must go to court to receive a judgment before it can have wages garnished or take the money from the defaulting borrower's bank account. Some states explicitly set the amount that the collection agency may impose to cover attorney fees, and when these are specified they are used in the construction of our variable. When actual amounts are not specified, but the collection of such fees is permissible within the state, we assign such fees the value of $50, which is the modal value among those states reporting what constitute "reasonable" attorney fees. In addition to the attorney fees, many states allow additional collection fees which are typically $10–$25, though a few states allow a much larger collection fee.

where $\boldsymbol{\psi}$ denotes the stacked vector of intercepts and covariate parameters in the system, $\boldsymbol{\eta}$ denotes the vector of component probabilities, $\boldsymbol{\delta}$ denotes the vector of cutpoints and IW denotes an inverse Wishart distribution. In practice we choose the hyperparameters of our priors as $\boldsymbol{\psi}_0 = \mathbf{0}_{k\times1}$, $\boldsymbol{V}_{\boldsymbol{\psi}} = 100\boldsymbol{I}_k$, $p = 8$, $\boldsymbol{R} = \boldsymbol{I}_3$ and $\kappa_1 = \kappa_2 = \cdots = \kappa_G = 1$, where k denotes the number of parameters contained in $\boldsymbol{\psi}$. These are rather conservative prior choices that have little influence on our posterior estimation results.

We close this section by offering some quick remarks regarding identification of the mixture components. While contention sometimes arises regarding this issue, we find this to be a minor concern for our analysis. First, we note that the slope coefficients are not component specific and thus no potential ambiguity arises regarding their interpretation. Second, our remaining objects of interest are a variety of posterior predictive densities and counterfactuals, as will be presented in Section 4, which are permutation invariant and thus, similarly unaffected by this problem.[9] Label switching, however, can be an important issue for model comparison and marginal likelihood calculation via simulated output; we briefly comment on this issue in the following section.

4. Results

We fit the model via the MCMC algorithm fully described in the Appendix, making use of data augmentation as described in Tanner and Wong (1987) and Albert and Chib (1993). We run the Gibbs algorithm for 100,000 iterations, discarding the first 20,000 draws as the burn-in period and retaining the last 80,000 simulations as draws from the joint posterior. Regression parameter posterior statistics will be presented in Table 4, with the fourth column of that table summarizing mixing performance through the calculation of numerical standard errors (NSEs) associated with the posterior mean estimates. In all cases, the reported NSEs are quite small relative to the posterior means, indicating that the number of simulations obtained under our algorithm enables reasonably precise estimation of those posterior moments.

Although we are encouraged by these results and the performance of our algorithm overall, our 80,000 draws clearly provide less information than the information that would be provided by an *i.i.d.* sample of equal size. For the slope coefficients on the exogenous variables, inefficiency factors are quite small (typically between 2 and 10), while the worst offenders proved to be the error correlation coefficients and coefficients on the endogenous variables. The inefficiency factors for these coefficients were often near or slightly in excess of 100, although our effective sample sizes in these worst cases still exceeded 700.

We considered models using one to four mixture components. Marginal likelihoods[10] suggested a strong preference for the two

component mixture model over the standard Gaussian model and also revealed a similarly strong preference for the use of two components over more highly parameterized three and four component specifications. This preference was also apparent when conducting diagnostic checking for the fit of each model. The standard single component Gaussian model was unable to reproduce key features of our data (such as skewness and bimodality in the loan amount distribution), whereas the two component model adequately reproduced those and other features. When expanding to a three component specification, we did not discern any noticeable improvement in overall model fit over that provided by the two component specification. For these reasons, we report results below from the two component mixture model only, and in the following section, also justify this choice through reasonably extensive model checking.

4.1. Assessing model fit

We investigate questions of model fit and adequacy by generating a series of simulations from the posterior predictive distribution, obtaining such simulations after first fixing all covariates to equal their in-sample values. The idea behind this exercise is to check a variety of dimensions of model fit by generating a series of outcomes from the maintained fitted model and then to see how well the distribution of those model-predicted outcomes align with the actual distribution of outcomes in our sample. Formally, we proceed by noting

$$p(\boldsymbol{y}^{\text{rep}}|\boldsymbol{y}^{\text{obs}}, \boldsymbol{x}^{\text{rep}} = \boldsymbol{x}^{\text{obs}})$$
$$= \int p(\boldsymbol{y}^{\text{rep}}|\boldsymbol{\theta}, \boldsymbol{y}^{\text{obs}}, \boldsymbol{x}^{\text{rep}} = \boldsymbol{x}^{\text{obs}})p(\boldsymbol{\theta}|\boldsymbol{y}^{\text{obs}}, \boldsymbol{x}^{\text{rep}} = \boldsymbol{x}^{\text{obs}}) \, d\boldsymbol{\theta}$$
$$= \int p(\boldsymbol{y}^{\text{rep}}|\boldsymbol{\theta}, \boldsymbol{x}^{\text{rep}} = \boldsymbol{x}^{\text{obs}})p(\boldsymbol{\theta}|\boldsymbol{y}^{\text{obs}}) \, d\boldsymbol{\theta}.$$

Therefore, for every post-convergence draw $\boldsymbol{\theta}^{(r)}$ from our posterior simulator, we generate a *replicated* data set from (1)–(8), denoted $\boldsymbol{y}^{\text{rep},(r)}$, upon setting the covariates $\boldsymbol{x}^{\text{rep}}$ to be equal to the covariates observed in-sample. Summaries of fit can then be obtained by comparing features of the replicated distribution to their counterparts obtained from the distribution of observed data.[11]

Fig. 1 presents one dimension of this exercise, as it plots a nonparametric estimate of the observed loan amount distribution alongside the posterior mean density estimate based upon the two component mixture version of (1)–(8). The latter is obtained by first noting that each post-convergence simulation yields a replicated loan amount vector. The estimated density of that vector is then obtained over a variety of fixed loan amount grid points. The resulting collection of density ordinates at each grid point – one obtained for each post-convergence simulation – is then averaged to obtain the plot in the figure.

In the right-most panel of Fig. 1, we also plot 90% HPD intervals, constructed from the collection of density ordinates at each grid point. As the left portion of the figure suggests, the estimated mixture density captures the skew and bimodality that is present in the loan amount data, and a standard single component Gaussian model clearly would not be sufficiently flexible to capture these features of the loan amount distribution. Finally, the right panel of Fig. 1 also illustrates that the nonparametric estimates fall within the mixture 90% HPD intervals at each grid point under consideration.

Another aspect of fitting the loan amount distribution concerns our ability to accurately predict how borrowers are constrained

---

[9] Geweke (2007) contains an excellent discussion of this and related issues.

[10] To calculate the marginal likelihoods, we follow the bridge sampling approach described in Frühwirth-Schnatter (2004), with related details supplied in Frühwirth-Schnatter (1994) and Frühwirth-Schnatter (2001), among others. This involved fitting the model without any component identification constraints and then randomly permuting the posterior simulations to achieve a "balanced" sampler. Details and code for this exercise are available upon request. Results (under equal prior odds) indicated that the two component model was favored over the single component Gaussian specification by a factor exceeding 6000:1. Adding more mixture components beyond the second was not supported by the data, as both the three and four component models had smaller marginal likelihoods than the single component Gaussian model. These results suggest the need to generalize normality, but not to proceed too far in that direction, as the added parameterization became unnecessary beyond two mixture components. Below, we document in detail the ability of the two component model to capture key features of our data.

[11] Lancaster describes this process in further detail Lancaster (2004, pp. 90–91) and eloquently summarizes it as "a re-run of history on the assumption that the model is what generates histories".

Table 3
Summary statistics for model fit.

(Renewal, default)	Full sample			Loans $\geq$ 300			Loans $\geq$ 500		
	Actual	$E(\cdot|\boldsymbol{y})$	Std$(\cdot|\boldsymbol{y})$	Actual	$E(\cdot|\boldsymbol{y})$	Std$(\cdot|\boldsymbol{y})$	Actual	$E(\cdot|\boldsymbol{y})$	Std$(\cdot|\boldsymbol{y})$
(0, 0)	0.145	0.141	0.011	0.059	0.045	0.007	0.013	0.012	0.003
(1, 0)	0.134	0.132	0.010	0.059	0.053	0.006	0.019	0.016	0.003
(2, 0)	0.107	0.105	0.009	0.053	0.047	0.006	0.020	0.016	0.003
(3, 0)	0.078	0.076	0.008	0.042	0.037	0.005	0.013	0.013	0.003
(4, 0)	0.058	0.057	0.007	0.032	0.029	0.004	0.013	0.010	0.002
(5, 0)	0.203	0.212	0.012	0.125	0.125	0.009	0.052	0.052	0.006
(0, 1)	0.155	0.158	0.011	0.037	0.032	0.005	0.011	0.009	0.002
(1, 1)	0.061	0.059	0.007	0.019	0.019	0.004	0.007	0.007	0.002
(2, 1)	0.029	0.028	0.005	0.009	0.010	0.003	0.004	0.004	0.002
(3, 1)	0.016	0.017	0.004	0.008	0.007	0.002	0.003	0.003	0.001
(4, 1)	0.012	0.012	0.003	0.005	0.006	0.002	0.003	0.003	0.001
(5, 1)	0.003	0.003	0.002	0.002	0.002	0.001	0.002	0.001	0.002

Fig. 1. Actual and replicated loan amount densities.

(or not constrained) by the state-imposed loan amount thresholds. Apart from the lender-imposed maximum of $1000 (which applies to states that have no explicit borrowing cap), the most common loan maxima in our data (as shown in Table 1) are $300 and $500. Among the states that impose a borrowing limit equal to $300, 42% of borrowers were observed to take out loans in the threshold amount. Our model matches this fraction rather closely as it predicts that 38% of such borrowers will seek a loan in the amount of $300, with an associated 90% HPD interval equal to (5%, 60%). Among the states that impose a borrowing limit equal to $500, 15% of borrowers are observed to take out a loan of the maximum amount. Our model predicts that 14.5% of borrowers in such states will seek to take out the maximum amount, with an associated 90% HPD interval equal to (0%, 24%). These results, together with those in Fig. 1, suggest that our model performs well in its ability to reproduce the distribution of loan amounts and in predicting the percentages of agents that are affected by state-imposed borrowing limits.

In terms of the number of loan renewals, the frequencies of renewal outcomes in our data are 30.0%, 19.5%, 13.6%, 9.4%, 7.0% and 20.6%, for *Renewals* $= 0, 1, 2, \ldots, 5$, respectively. Not surprisingly, our ordinal choice model matches these outcomes very closely, as the posterior mean frequencies are 29.9%, 19.1%, 13.3%, 9.3%, 6.9% and 21.5%, respectively. The overall probability of default is estimated equally well, with the observed frequency of default and its posterior mean both found to equal 28%.

Table 3 expands on this replication exercise and further tests the performance of our model. In the first block of columns of

Table 3, we report results associated with replication of the full joint distribution of renewal and default outcomes. Columns 2–3 of the table report the observed (Actual) and estimated $[E(\cdot|\boldsymbol{y})]$ *joint* probabilities of each renewal/default cell. The third column reports the posterior standard deviation associated with each cell probability. Here we see that our model is able to replicate this finer set of discrete probabilities very well, as the posterior means match the actual frequencies closely in all cases. Furthermore, the posterior standard deviations associated with these fractions are also reasonably small, indicating the cell probabilities are precisely estimated by our model.

The remaining sets of columns of Table 3 examine how well the model is able to reproduce the joint distribution of renewal and default outcomes across different loan amount values. In the middle portion of the table (columns 5–7), we report observed and estimated joint probabilities of renewal/default pairs *and* the event that the loan will be at least as large as $300. A similar exercise is conducted in columns 8–10 of the table, this time presenting various joint probabilities associated with loans at least as large as $500. Again, we see that the observed and estimated probabilities align rather closely, indicating that our model is able to reproduce the observed renewal/default joint distribution not just unconditionally, but also across different loan sizes. In these last two replication exercises, our point estimates of the cell probabilities fall within one posterior standard deviation of the observed outcome in 21 out of 24 cases. In the remaining three instances, the observed outcomes fall within two standard deviations of the posterior means.

4.2. Parameter estimates and marginal impacts

Regression parameter posterior means, posterior standard deviations, numerical standard errors and probabilities of being positive associated with the two component mixture model are presented in Table 4.

Given the nonlinearity of the model, the coefficients themselves are obviously not directly interpretable as marginal effects. For this reason, we report in the fifth column of Table 4 a series of point estimates of the marginal impacts (denoted "ME") of the different covariates on each equation's outcome.[12] For a given post-convergence draw from our simulator, we calculate how much the outcome of interest changes following a particular change in the given covariate, and do this for every observation

---

[12] The marginal effects in this table are not complete summaries as they do not capture any "indirect effects" of the covariate that feed through the model. For example, in calculating the effect of *LogRent* on *Renewals*, we do not consider the effect of *LogRent* on the amount borrowed and the subsequent impact of the amount borrowed on renewal behavior. Complete summaries of this type are, however, considered in the following section.

Table 4
Posterior summary statistics of a selection of parameters (from β_{Rate} to δ_5).

Parameter (ξ)	$E(\xi\|D)$	$\text{Std}(\xi\|D)$	$\text{Pr}(\xi > 0\|D)$	NSE	ME
β_{Rate}	−0.0483	0.298	0.438	0.00165	−1.18
β_{Term}	0.0021	0.00224	0.825	1.29e−005	0.536
$\beta_{LogRent}$	0.0481	0.0286	0.954	0.000166	1.39
$\beta_{LogStateMAX}$	0.0828	0.03	0.997	0.000172	11.4
γ_{LogAmt}	0.518	0.128	1	0.0024	0.405
γ_{Rate}	−1.06	0.641	0.0497	0.00626	−0.234
γ_{Term}	0.0104	0.00341	0.999	1.48e−005	0.0236
$\gamma_{LogRent}$	−0.0166	0.0435	0.351	0.000187	−0.00423
$\gamma_{StateCap}$	0.0343	0.0134	0.994	0.00019	0.0523
$\gamma_{NoStateCap}$	0.131	0.0551	0.991	0.00069	0.298
θ_{LogAmt}	0.693	0.303	0.991	0.0125	0.05
$\theta_{I(Renewals=1)}$	−0.38	0.179	0.0145	0.00606	−0.0287
$\theta_{I(Renewals=2)}$	−0.64	0.246	0.00356	0.0091	−0.0575
$\theta_{I(Renewals=3)}$	−0.683	0.294	0.00835	0.0109	−0.0632
$\theta_{I(Renewals=4)}$	−0.633	0.337	0.0292	0.0127	−0.0552
$\theta_{I(Renewals=5)}$	−1.98	0.465	0	0.0199	−0.21
θ_{Rate}	2.25	0.99	0.987	0.00947	0.0503
θ_{Term}	0.028	0.00891	1	0.000212	0.0059
$\theta_{LogRent}$	−0.447	0.112	0	0.00189	−0.0107
$\theta_{LogStatePenalty}$	−0.181	0.055	0.000175	0.0014	−0.0373
$\theta_{StatuteLimit}$	−0.0176	0.0132	0.0886	0.000109	−0.00369
δ_3	0.384	0.0142	1	0.000262	
δ_4	0.641	0.0145	1	0.000279	
δ_5	0.835	0.0116	1	0.000219	

in the sample. The results are then averaged over the collection of sample values. The process is repeated for all posterior simulations and then averaged a final time to produce the entries in the last column of the table.

Marginal effects related to *LogRent*, *LogAmt*, *LogStateMax* and *LogStatePenalty* are calculated by considering changes in the level of those variables equal to $100, while marginal effects related to *Rate* estimate impacts resulting from increasing *Rate* by 0.1. Unless otherwise noted, marginal effects for the remaining parameters are associated with unit increases in the variable under consideration. Estimated MEs for the *LogAmt* equation represent expected changes in the amount borrowed (not its log). For the default equation, the MEs quantify the increase or decrease in the probability of default while effects for the *Renewals* equation simply report the expected increase in the number of times the loan is renewed.

We first observe from Table 4 that the interest rate associated with the loan has little impact on loan size [$\text{Pr}(\beta_{Rate} > 0|\boldsymbol{y}) = 0.438$, loan amounts fall by only $1.18 on average in response to increasing *Rate* by 0.1], which is consistent with our earlier observation that a majority of payday borrowers take out loans for emergency purposes and may have no other recourse for managing an unanticipated economic shock. The interest rate does, however, have a rather strong effect on the number of loan renewals [$\text{Pr}(\gamma_{Rate} > 0|\boldsymbol{y}) \approx 0.05$], and increasing the rate by 0.1 lowers the expected number of renewals by 0.23. Higher interest rates also result in increased likelihoods of default [$\text{Pr}(\theta_{Rate} > 0|\boldsymbol{y}) \approx 0.99$], and increasing the interest rate by 0.1 increases the probability of default by approximately 5%. Given these results, the expected impact of an interest rate reduction on lender profits is unclear and warrants additional investigation: although a higher interest rate allows the lender to collect higher fees at each renewal, such an effect may be partially or wholly offset by fewer interest payments in total and more loans ending in default. We will return to this issue when discussing several different policy simulations in the following section.[13]

The (log) monthly rent paid by the individual is also found to be an important determinant of loan size and the decision to default. Borrowers who pay more in rent each month tend to take out larger loans [$\text{Pr}(\beta_{LogRent} > 0|\boldsymbol{y}) > 0.95$]. Specifically, a $100 increase in monthly rent paid is associated with an expected increase in loan amount equal to $1.39, a rather small quantity.[14] Individuals paying more per month in rent are also less likely to default [$\text{Pr}(\theta_{LogRent} > 0|\boldsymbol{y}) \approx 0$], with borrowers paying $100 more per month in rent being approximately 1.1% less likely to default on their loans. Rent, however, is found to have little economic [increasing rent by $100 decreases the number of renewals by 0.004] or statistical [$\text{Pr}(\gamma_{LogRent} > 0|\boldsymbol{y}) \approx 0.35$] impact on the number of loan renewals. This pattern of outcomes is consistent with what we would expect, given that individuals with greater financial stability (and thus presumably paying more in rent) should be associated with larger loans and fewer incidences of default.

Larger loans are clearly renewed more often, as $\text{Pr}(\gamma_{LogAmt} > 0|\boldsymbol{y}) \approx 1$ (that is, all posterior simulations associated with this parameter were positive). Increasing loan size by $100 leads to an expected increase in the number of renewals by approximately 0.41. Furthermore, and perhaps not surprisingly, larger loans are more likely to end in default, as $\text{Pr}(\theta_{LogAmt} > 0|\boldsymbol{y}) \approx 0.99$. Holding all other factors constant, increasing the loan amount by $100 results in a 5% increase in the probability of default.

The state-level policy variables employed as exclusion restrictions all seem to play prominent roles and operate in the directions

---

[13] We do not model the extensive margin aspect of this problem, by which we mean that lower rates might potentially induce more individuals to seek out payday loans who otherwise would not. We focus, instead, solely on the "intensive margin", given the data available to us. In a similar vein, one might be concerned that a policy

change could lead to changes in the types of individuals that take out payday loans, which is not fully accounted for in our analysis. While this is certainly possible, and our results must be interpreted with a degree of caution in light of this, we do not believe this is a significant problem. On the lender side, most loans are simply approved and thus little supply-side selection takes place. On the demand side, we do not find strong relationships between state policy parameters (such as the interest rate charged, maximum number of renewals and default penalties) and either the rents paid by our borrowers or zipcode-level characteristics (such as median income and education levels) associated with those borrowers.

[14] Monthly rents in our data range from $367 to $2388, with a standard deviation in excess of $350. Thus, a $100 change in rent represents a rather modest increase in the level of that variable. Furthermore, the desired increase in loan amount associated with the rent increase is not permitted by prevailing state policy in many cases, as the latent amount is already in excess of the state borrowing limit. Marginal effects in such instances are necessarily coded as zero and thus attenuate the overall impact.

we expect. In addition to effects operating through support restrictions, states having higher borrowing maxima are also associated with higher conditional means in the latent loan amount equation, as $\Pr(\beta_{LogStateMax} > 0|\boldsymbol{y}) \approx 1$. Increasing the borrowing limit by $100 results in loans that are $11.40 larger on average. States with fewer restrictions on the number of permissible renewals are associated with loans being held longer [$\Pr(\gamma_{StateCap} > 0|\boldsymbol{y}) \approx 1$, $\Pr(\gamma_{NostateCap} > 0|\boldsymbol{y}) \approx 1$] and borrowers residing in states with more stringent penalties on defaulters are less likely to default: $\Pr(\theta_{LogStatePenalty} > 0|\boldsymbol{y}) \approx 0$. Similarly, states that allow more time for the lender to collect on negligent debt are generally associated with lower probabilities of default [$\Pr(\theta_{StatuteLimit} > 0|\boldsymbol{y}) > 0 \approx 0.089$].

Point estimates of the loan renewal indicators in the default equation generally reveal a monotonic pattern,[15] and indicate that those renewing the loan at least once are significantly less likely to default than those who never renew. This finding is, perhaps, not terribly surprising given the nearly even split of 0 renewal outcomes among defaulters and non-defaulters and an overall sample default rate of 28%. The coefficients also reveal that those renewing the loan five times – the cap imposed by the lender – are very unlikely to default, and this group is clearly differentiated from those renewing between one and four times. The explanation here is that agents realize this is the last renewal possibility and thus choose to renew a final time only if they know they will be able to pay off the loan in full.[16] The correlation between the errors of the renewals and default equations was also found to be negative, though rather imprecisely estimated. This negative correlation is consistent with the view that agents who renew a loan more often than expected (and thus have relatively high u values) are precisely the ones that have every intention of paying off the loan and do everything they can to do so (and thus tend to have relatively low values of v).

4.3. Model predictions and counterfactual exercises

The previous discussion of results read mostly as an accounting exercise, where we simply enumerated variables that were empirically important, discussed the directions of their effects, and in some cases offered brief discussions of particular marginal impacts. A primary goal of this paper, however, is more ambitious than this, as we seek to use our model to describe how changes in policy variables—such as the interest rate and state loan maximum—will impact loan amounts, loan durations, default decisions and lender revenues. Importantly, changes in these policy variables filter through the entire equation system and affect all our outcomes through a variety of direct and indirect channels. In what follows we thus determine the total impacts of various policy changes on each of our three endogenous variables and also investigate how lender revenues are affected under different policy experiments.

To this end, we again look to the posterior predictive distribution. Let

$$\boldsymbol{y}_f = [logAmt_f \ Renewals_f \ Default_f]$$

and define $\boldsymbol{y}_f^*$ similarly as future unobserved outcomes that we wish to characterize given the model. Dropping superfluous conditioning information as appropriate, the posterior predictive distribution for this vector of outcomes is obtained as:

$$p(\boldsymbol{y}_f|\boldsymbol{y}, \boldsymbol{x}_f) = \int_{\Theta} p(\boldsymbol{y}_f|\boldsymbol{y}_f^*)p(\boldsymbol{y}_f^*|\boldsymbol{\theta}, \boldsymbol{x}_f, \boldsymbol{y})p(\boldsymbol{\theta}|\boldsymbol{y})d\boldsymbol{\theta}. \quad (13)$$

Samples from (13) can be easily generated via the method of composition, given the marginal-conditional decomposition of the integrand.

We use this basic idea to perform three different counterfactual exercises.[17] For the first of these, we consider changes in loan amount, loan duration and default that result from a state lowering its maximum interest rate to 0.10 (the lowest value in our sample) from an initial rate of 0.23 (the highest value in the sample). Second, we consider the change in these same variables as a result of a state lowering the maximum amount that can be borrowed from $1000 (the largest value in the sample) to $300 (the smallest sample value). Third and finally, we consider what would happen if a state moved from imposing no cap on the maximum number of renewals a borrower could make to imposing (although not actually being able to enforce) that a borrower cannot renew the loan and must pay in full at the next payday.

Results of these policy experiments are provided in Table 5 and some explanations of the quantities provided in that table and their calculation are in order. As the notation of (13) makes clear, draws from (13) are conditional upon values of the covariates $\boldsymbol{x}_f$. To eliminate this dependence in our results, we follow Chib and Jacobi (2007) and sample values of the covariates that are not involved in the construction of our counterfactual from the empirical distribution of covariates in our sample. We do this for each iteration of the sampler. This enables us to report average (and unconditional) posterior predictive impacts of the policy changes instead of simply reporting impacts for a particular subset of individuals.

The first column of Table 5 summarizes effects of the policy changes on the amount borrowed. The second column summarizes the effects on the number of times that the loan is renewed. Results in this second column offer a complete summary of the impact of the policy change on loan renewals in the sense that they take into account the "direct" effect of the change in policy variable as well as its indirect effect derived from changes to the loan amount. The third column summarizes the effects on the default outcome. Again, these offer a complete summary on the default impact as they account for both the direct effect of the policy change as well as indirect effects that filter through from changes to loan amount and to loan duration.

The final column of Table 5 describes the effect on revenues to the lender. For a given draw from (13), the amount returned to the lender from a loan provided under the current set of policy variables is directly calculable as[18]:

$$LenderRevenue_f = Default_f(\exp(logAmt_f) * Renewals_f * Rate_f)$$
$$+ (1 - Default_f) * (\exp(logAmt_f)$$
$$\times [1 + (1 + Renewals_f) * Rate_f]). \quad (14)$$

---

[15] An exception to this occurs when moving from three to four renewals.

[16] An agent faced with a fifth renewal decision who realizes he/she cannot repay the loan would rationally default after the fourth renewal decision and choose not to renew at this final opportunity. Such backward induction is not applied in all periods, however, as it is clearly optimal for defaulting agents to do so immediately rather than make one or several interest payments only to eventually default. Since the time from loan receipt to final repayment can be quite long—nearly 3 months for biweekly paid individuals and 6 months for monthly paid individuals—borrowers may face considerable uncertainty in their income streams over this period and thus continue to renew with the expectation that repayment will eventually be possible. For individuals facing the final renewal decision, this type of uncertainty is lessened, and those renewing the maximum amount typically do not default.

[17] For these calculations we keep the unobservables constant across the two different scenarios as we envision tracing out the effects of these policy changes for a given individual.

[18] This calculation does not explicitly account for a few aspects of the payday lending process. First, it ignores the salvage value of a defaulted loan that is recoverable by the lender upon selling such loans to a collection agency. Second, a very small number of states in our data permit the lender to charge an initial fee for processing the loan (typically less than $10) as well as similar fees upon each loan renewal. As these fees are quite uncommon in our data, we do not consider them when calculating (14). Third, we also assume that, at each loan renewal, the agent simply pays the interest and does not pay down any principal. Finally, we do not consider the case where loans are repaid within 3 or 4 days since such cases of rapid repayment cannot be charged any interest by the lender.

Table 5
Counterfactual calculations.

Statistic	Amount	Renewals	Default	Revenue
Interest rate				
$E(\Delta|\boldsymbol{y})$	$1.69	0.324	−0.083	$−68.7
$Std[\Delta|\boldsymbol{y}]$	$16.99	0.513	0.280	$162.8
$Pr(\Delta > 0|\boldsymbol{y})$	0.51	0.310	0.002	0.121
$Pr(\Delta = 0|\boldsymbol{y})$	0.10	0.686	0.913	0.127
$Pr(\Delta < 0|\boldsymbol{y})$	0.39	0.004	0.085	0.752
State loan maximum				
$E(\Delta|\boldsymbol{y})$	$−100.50	−0.325	−0.058	$−114.78
$Std[\Delta|\boldsymbol{y}]$	$152.65	0.624	0.235	$245.65
$Pr(\Delta > 0|\boldsymbol{y})$	0.002	0.000	0.000	0.048
$Pr(\Delta = 0|\boldsymbol{y})$	0.000	0.745	0.942	0.144
$Pr(\Delta < 0|\boldsymbol{y})$	0.998	0.255	0.058	0.808
State renewal cap				
$E(\Delta|\boldsymbol{y})$	–	−0.298	0.025	$−62.87
$Std[\Delta|\boldsymbol{y}]$	–	0.474	0.173	$197.99
$Pr(\Delta > 0|\boldsymbol{y})$	–	0.028	0.028	0.003
$Pr(\Delta = 0|\boldsymbol{y})$	–	0.745	0.970	0.709
$Pr(\Delta < 0|\boldsymbol{y})$	–	0.255	0.002	0.288

The first five rows of Table 5 present results from our interest rate reduction experiment. Here, we again see that the rate drop has a rather small impact on the size of the loan. The expected increase in loan size resulting from lowering the interest rate from 0.23 to 0.10 is merely $1.69, with a posterior standard deviation estimated to be ten times as large as the posterior mean. Approximately ten percent of the time the amount of the loan is completely unaffected by lowering the interest rate in this way, as these outcomes occur when the borrower seeks to take out the state maximum amount under both interest rate scenarios.[19] The interest rate drop also leads to an expected increase in the number of loan renewals equal to 0.32 and increases the number of loan renewals approximately 31% of the time. Finally, the interest rate reduction lowers the probability of default by approximately 8.3% while the posterior probability that default behavior is completely unaffected is approximately 0.91.

Although the interest rate reduction leads to (very) slightly larger loans on average, loans that are renewed more often, and an overall decrease in the probability of default, the mean effect of this reduction on lender revenues remains negative, lowering the average amount returned to the lender by $68.7. Seventy-five percent of the time lender revenues decline as a result of this rate reduction, and revenues are completely unaffected 13% of the time. These instances of no revenue change occur when either the borrower takes out the state maximum under both rates and renews the loan the same number of times or when the agent defaults without any form of repayment under both regimes. Lender revenues are observed to increase with the interest rate reduction in approximately 12% of loans. These positive outcomes always occur when agents are induced to repay their loan under the lower rate (but otherwise would have defaulted), while other positive revenue changes are realized when the agent is induced to renew the loan additional times when facing a lower interest rate. Among the revenue difference simulations that are "large" – defined here in excess of $900 in absolute value – 92% of these

correspond to revenue gains under the lower rate regime. These realizations correspond exactly to the cases where default behavior is affected—having a loan successfully paid in full an additional 8.3% of the time can result in huge revenue gains in the low rate regime relative to the high rate regime. Such concerns over default behavior and the large losses that can occur with default likely motivated our lender's decision to set 0.23 as the charged interest rate even in states that allowed higher rates or provided no explicit rate cap.

The second portion of the table provides posterior statistics associated with dropping the state loan maximum from $1000 to $300. Here the results are clear and mostly unfavorable from the lender's perspective. Lowering the state maximum in this fashion results in loans that are smaller by approximately $100 on average and, not surprisingly, more restrictive borrowing limits almost always lower the size of the loan.[20] The reduction in the loan cap also lowers the expected number of loan renewals by 0.33, is never associated with additional renewals in any of our simulations, and results in fewer loan renewals approximately 26% of the time. Lowering the state cap never induces new cases of default in any of our simulations and eliminates default in approximately 5.8% of loans. The impacts on renewals and default, of course, are primarily driven by the smaller loan sizes themselves, which were found in Table 4 to reduce the number of loan renewals and to discourage default. Finally, lender revenue is reduced by approximately $115 on average when lowering the state maximum to $300, and revenues are lowered nearly 81% of the time. In 14% of cases, lender revenues are unaffected when restricting the maximum loan amount in this way, and these cases correspond to immediate defaults by the borrower regardless of the permissible loan size. Lender revenues also increase approximately 5% of the time, all of which arise when the lower loan size induces the agent to pay off the loan rather than default.

The third and final set of rows in Table 5 present results related to the state renewal cap counterfactual. In this case we investigate what happens to our outcome variables of interest when a state moves from a policy of no restrictions on the number of renewals taken to one where the borrower is restricted to pay off the loan in full at the next pay period. Given the assumed triangularity of the model, this change has no impact on the size of the loan, but does reduce the expected number of renewals by approximately 0.3. The imposition of the renewal cap has a small impact on default behavior, leaving default outcomes unchanged 97% of the time and inducing default in approximately 3% of loans. The increased likelihood of default results from the overall reduction in loan renewals, as Table 4 revealed a general trend that loans renewed more often were less likely to end in default. Lender revenues decrease with probability 0.29 and in an average amount of $62.9 as a result of prohibiting loan renewals, owing to both fewer interest payments made on average and a small increase in the default likelihood. This situation, somewhat unlike the other two experiments discussed earlier, reveals little opportunity for an increase in lender revenues as a result of the policy change, as the posterior probability that revenues will increase following the ban on renewals is just 0.003.

As a whole these results suggest that state-level regulations have important effects on borrower behavior. If policy makers wish to mitigate default, reducing the maximum interest rate that can be charged and the maximum amount that can be borrowed offer somewhat effective ways to lessen its incidence. On the other hand, prohibiting loan renewals has comparably little impact on

---

[19] For example, consider that $Pr(\beta_{Rate} > 0|\boldsymbol{y}) = 0.438$ in Table 4. Positive values of the rate coefficient would seem to be associated with *negative* values of Δ_{Amount} in the first column of Table 5, as the effect we consider pertains to a sizeable interest rate *reduction*. However, we note $Pr(\Delta_{Amount} < 0|\boldsymbol{y}) = 0.39$ in Table 5, a smaller number than 0.438. This slight discrepancy arises since some positive β_{Rate} simulations produce exactly zero change in the loan amount. This happens when the borrower wishes to take out a loan at least as large as the maximum allowable amount even under the low rate regime and thus the change to amount borrowed is identically zero.

---

[20] From Table 4 we see that $\beta_{LogStateMax}$ is negative for 0.3% of the simulations, and these draws were associated with larger loan amounts or loan amounts that were unchanged under both scenarios.

default behavior and restricting the payoff deadline in this way may actually induce some individuals to default.

Our revenue calculations also suggest that payday lenders would more strongly oppose legislation that restricts the maximum loan amount than legislation that would reduce the maximum interest rate (at least over the supports spanned by our data). In fact, sufficiently risk-averse lenders could potentially prefer a lower interest rate given its modest preventative effects on default behavior and the large gains that are realized when lower rates induce agents to pay off their loans. Such a politically popular effort to lower interest rates would be met with less (or perhaps little) resistance by payday lenders if it were accompanied by more stringent penalties on defaulting borrowers. The results of Table 4, for example, show such penalties to be rather effective, as a $100 increase in the amount that can be recovered on a defaulted $300 loan lowers the default rate by nearly 4%.

5. Conclusion

Payday lending is a common practice with an estimated ten million borrowers in the US (Skiba and Tobacman, 2008). Researchers have primarily addressed the question of whether access to payday lending is beneficial or harmful to borrowers and the results have been mixed. While payday lending is illegal in some states, most states allow payday lending and regulate the terms of the loans. There is great interest by policy makers, consumer-rights advocates, and the payday loan industry to know how borrowers respond to these state-level regulations.

Payday loan regulations vary substantially across states, and we make use of data from an online lender with loans distributed across 38 states. We use flexible Bayesian methods to estimate a model that allows us to describe how these state-level regulations affect borrower behavior. To our knowledge, this is the first paper to investigate these relationships.

Our estimates suggest that decreasing the maximum interest rate that may be charged increases the length of time the loan is held and decreases the probability of default. Reducing the maximum amount that an individual may borrow decreases the amount individuals choose to borrow (even those not constrained by the maximum), decreases the length of time the loan is held, and also decreases the probability of default. Requiring borrowers to repay their entire loan on their next payday (rather than allowing loan renewals to continue indefinitely) results in fewer loan renewals on average, lower lender revenues and an approximate 3% increase in the likelihood of default.

Appendix

In this section we provide details of our posterior simulator. The model in (1)–(8) is a linear SEM on suitably defined latent data. Posterior simulation proceeds via MCMC methods, as described completely below. In the description that follows we assume the number of mixture components, G, is given.

To simplify some of our later expressions, we first define:

$a_i^* = LogAmt_i^*$

$\boldsymbol{x}_i = (Rate_i, Term_i, LogRent_i, LogStateMAX_i)$

$\boldsymbol{\beta} = (\beta_1, \beta_2, \beta_3, \beta_4)'$

$a_i = LogAmt_i$

$m_i = LogStateMAX_i$

$r_i^* = Renewals_i^*$

$\boldsymbol{z}_i = (LogAmt_i, Rate_i, Term_i, LogRent_i, StateCap_i, NoStateCap_i)$

$\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6)'$

$r_i = Renewals_i$

$\boldsymbol{\delta} = (\delta_1 = -\infty, \delta_2 = 0, \delta_3, \delta_4, \delta_5, \delta_6 = 1, \delta_7 = \infty)$

$d_i^* = Default_i^*$

$\boldsymbol{w}_i = (LogAmt_i, I(Renewals_i = 1), \dots, I(Renewals_i = 5),$
$\qquad Rate_i, Term_i, LogRent_i, LogStatePenalty_i, StatuteLimit_i)$

$\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_{10}, \theta_{11})'$

$d_i = Default_i.$

With this notation in hand, we can write our model as, conditional on the value of the component labeling variable c_i:

$a_i^* = \beta_{0c_i} + \boldsymbol{x}_i \boldsymbol{\beta} + \epsilon_i \equiv \tilde{\boldsymbol{x}}_i \tilde{\boldsymbol{\beta}} + \epsilon_i$

$a_i = \min\{a_i^*, m_i\}$

$r_i^* = \gamma_{0c_i} + \boldsymbol{z}_i \boldsymbol{\gamma} + u_i \equiv \tilde{\boldsymbol{z}}_i \tilde{\boldsymbol{\gamma}} + u_i$

$r_i = \sum_{j=0}^{5} j \times I(\delta_{j+1} < r_i^* \le \delta_{j+2})$

$d_i^* = \delta_{0c_i} + \boldsymbol{w}_i \boldsymbol{\theta} + v_i \equiv \tilde{\boldsymbol{w}}_i \tilde{\boldsymbol{\theta}} + v_i$

$d_i = I(d_i^* > 0)$

$(\epsilon_i, u_i, v_i)' | \cdot, \boldsymbol{c}_i \overset{\text{ind}}{\sim} N(\boldsymbol{0}_{3 \times 1}, \boldsymbol{\Sigma}_{c_i})$

where $\tilde{\boldsymbol{x}}_i = \tilde{\boldsymbol{x}}_i(c_i) \equiv [I(c_i = 1) \, I(c_i = 2) \cdots I(c_i = G) \, \boldsymbol{x}_i]$, $\tilde{\boldsymbol{\beta}} = [\beta_{01} \beta_{02} \cdots \beta_{0G} \, \boldsymbol{\beta}']'$ and $\tilde{\boldsymbol{z}}_i, \tilde{\boldsymbol{w}}_i, \tilde{\boldsymbol{\gamma}}$ and $\tilde{\boldsymbol{\theta}}$ are defined similarly.

Let $\boldsymbol{d}^* = (d_1^*, d_2^*, \dots, d_n^*)'$, $\boldsymbol{d} = (d_1, d_2, \dots, d_n)'$, $\boldsymbol{r}^* = (r_1^*, r_2^*, \dots, r_n^*)'$, $\boldsymbol{r} = (r_1, r_2, \dots, r_n)'$, $\boldsymbol{a}^* = (a_1^*, a_2^*, \dots, a_n^*)'$, $\boldsymbol{a} = (a_1, a_2, \dots, a_n)'$, $\boldsymbol{c} = (c_1, c_2, \dots, c_n)'$, $\boldsymbol{\Sigma} = (\boldsymbol{\Sigma}_1, \boldsymbol{\Sigma}_2, \dots, \boldsymbol{\Sigma}_G)$, and $\boldsymbol{\psi} = (\tilde{\boldsymbol{\beta}}', \tilde{\boldsymbol{\gamma}}', \tilde{\boldsymbol{\theta}}')'$. The priors employed were given in (9)–(12).

The augmented joint posterior distribution of $\boldsymbol{d}^*$, $\boldsymbol{r}^*$, $\boldsymbol{a}^*$, $\boldsymbol{c}$, $\boldsymbol{\psi}$, $\boldsymbol{\eta}$, $\boldsymbol{\delta}$ and $\boldsymbol{\Sigma}$ is thus proportional to the following joint density of observed data, latent data and model parameters:

$$p(\boldsymbol{d}^*, \boldsymbol{d}, \boldsymbol{r}^*, \boldsymbol{r}, \boldsymbol{a}^*, \boldsymbol{a}, \boldsymbol{c}, \boldsymbol{\eta}, \boldsymbol{\psi}, \boldsymbol{\delta}, \boldsymbol{\Sigma})$$

$$\propto \left\{ \prod_{g=1}^{G} \eta_g^{\kappa_g - 1} |\boldsymbol{\Sigma}_g|^{-\frac{p+4}{2}} \exp\left[-\frac{1}{2} \text{tr}(p\boldsymbol{R}\boldsymbol{\Sigma}_g^{-1}) \right] I(\sigma_{vvg} = 1) \right\}$$

$$\times |2\pi \boldsymbol{V}_{\boldsymbol{\psi}}|^{-\frac{1}{2}}$$

$$\times \exp\left[-\frac{1}{2} (\boldsymbol{\psi} - \boldsymbol{\psi_0})' \boldsymbol{V}_{\boldsymbol{\psi}}^{-1} (\boldsymbol{\psi} - \boldsymbol{\psi_0}) \right] I(\delta_1 = -\infty)$$

$$\times I(\delta_2 = 0) I(0 < \delta_3 < \delta_4 < \delta_5 < 1) I(\delta_6 = 1) I(\delta_7 = \infty)$$

$$\times \prod_{i=1}^{n} |2\pi \boldsymbol{\Sigma}_{c_i}|^{-\frac{1}{2}}$$

$$\times \exp\left[-\frac{1}{2} \begin{pmatrix} a_i^* - \tilde{\boldsymbol{x}}_i \tilde{\boldsymbol{\beta}} \\ r_i^* - \tilde{\boldsymbol{z}}_i \tilde{\boldsymbol{\gamma}} \\ d_i^* - \tilde{\boldsymbol{w}}_i \tilde{\boldsymbol{\theta}} \end{pmatrix}' \boldsymbol{\Sigma}_{c_i}^{-1} \begin{pmatrix} a_i^* - \tilde{\boldsymbol{x}}_i \tilde{\boldsymbol{\beta}} \\ r_i^* - \tilde{\boldsymbol{z}}_i \tilde{\boldsymbol{\gamma}} \\ d_i^* - \tilde{\boldsymbol{w}}_i \tilde{\boldsymbol{\theta}} \end{pmatrix} \right]$$

$$\times [I(a_i = a_i^*) \times I(a_i^* < m_i) + I(a_i = m_i) \times I(a_i^* \ge m_i)]$$

$$\times I(\delta_{r_i+1} < r_i^* \le \delta_{r_i+2}) [I(d_i = 0)$$

$$\times I(d_i^* \le 0) + I(d_i = 1) \times I(d_i^* > 0)]$$

$$\times \left[\sum_{g=1}^{G} \eta_g I(c_i = g) \right]. \tag{15}$$

We generate samples from the posterior distribution derived from (15) via a combination of Gibbs and Metropolis–Hastings steps. Each iteration involves a total of eight sampling steps which are enumerated in detail below:

Step 1: Sampling $\boldsymbol{c}$

For $i = 1, 2, \dots, n$, the component label variable c_i is sampled from a discrete distribution where the conditional posterior

probability that $c_i = g$, for $g = 1, 2, \ldots, G$, is simply

$\Pr(c_i = g | \boldsymbol{\Xi}_{-c_i}, \boldsymbol{D})$

$$= \frac{\eta_g |\boldsymbol{\Sigma}_g|^{-\frac{1}{2}} \exp\left[-\frac{1}{2} \begin{pmatrix} a_i^* - \beta_{0g} - \boldsymbol{x}_i\boldsymbol{\beta} \\ r_i^* - \gamma_{0g} - \boldsymbol{z}_i\boldsymbol{\gamma} \\ d_i^* - \theta_{0g} - \boldsymbol{w}_i\boldsymbol{\theta} \end{pmatrix}' \boldsymbol{\Sigma}_g^{-1} \begin{pmatrix} a_i^* - \beta_{0g} - \boldsymbol{x}_i\boldsymbol{\beta} \\ r_i^* - \gamma_{0g} - \boldsymbol{z}_i\boldsymbol{\gamma} \\ d_i^* - \theta_{0g} - \boldsymbol{w}_i\boldsymbol{\theta} \end{pmatrix} \right]}{\sum\limits_{h=1}^{G} \eta_h |\boldsymbol{\Sigma}_h|^{-\frac{1}{2}} \exp\left[-\frac{1}{2} \begin{pmatrix} a_i^* - \beta_{0h} - \boldsymbol{x}_i\boldsymbol{\beta} \\ r_i^* - \gamma_{0h} - \boldsymbol{z}_i\boldsymbol{\gamma} \\ d_i^* - \theta_{0h} - \boldsymbol{w}_i\boldsymbol{\theta} \end{pmatrix}' \boldsymbol{\Sigma}_h^{-1} \begin{pmatrix} a_i^* - \beta_{0h} - \boldsymbol{x}_i\boldsymbol{\beta} \\ r_i^* - \gamma_{0h} - \boldsymbol{z}_i\boldsymbol{\gamma} \\ d_i^* - \theta_{0h} - \boldsymbol{w}_i\boldsymbol{\theta} \end{pmatrix} \right]},$$

$\forall i, g$.

Step 2: Sampling $\boldsymbol{\eta}$

From (15), it is seen that the posterior conditional distribution of the component probability vector $\boldsymbol{\eta}$ is Dirichlet:

$\boldsymbol{\eta} | \boldsymbol{\Xi}_{-\eta}, \boldsymbol{D} \sim \text{Dirichlet}(\kappa_1 + n_1, \kappa_2 + n_2, \ldots, \kappa_G + n_G)$

where $n_g \equiv \sum_{i=1}^n I(c_i = g)$, $g = 1, 2, \ldots, G$.

Step 3: Sampling $\boldsymbol{\Sigma}_g$

The conditional posterior distributions of the component-specific covariance matrices $\boldsymbol{\Sigma}_g$ are independent, with each proportional to

$p(\boldsymbol{\Sigma}_g | \boldsymbol{\Xi}_{-\Sigma_g}, \boldsymbol{D})$

$$\propto p_{IW}\left[\boldsymbol{\Sigma}_g \,\middle|\, p + n_g, p\boldsymbol{R} + \sum_{\{i:c_i=g\}} \right.$$
$$\left. \times \begin{pmatrix} a_i^* - \beta_{0g} - \boldsymbol{x}_i\boldsymbol{\beta} \\ r_i^* - \gamma_{0g} - \boldsymbol{z}_i\boldsymbol{\gamma} \\ d_i^* - \theta_{0g} - \boldsymbol{w}_i\boldsymbol{\theta} \end{pmatrix} \begin{pmatrix} a_i^* - \beta_{0g} - \boldsymbol{x}_i\boldsymbol{\beta} \\ r_i^* - \gamma_{0g} - \boldsymbol{z}_i\boldsymbol{\gamma} \\ d_i^* - \theta_{0g} - \boldsymbol{w}_i\boldsymbol{\theta} \end{pmatrix}' \right] I(\sigma_{vvg} = 1),$$

$g = 1, 2, \ldots, G$.

To generate draws from this conditional posterior, we make use of Nobile's (2000) comments on McCulloch et al.'s (2000) reparameterization approach to posterior simulation in multinomial probit (MNP) models. The problem we face in step 3 is identical to that faced in the MNP model, as it demands generating draws from an inverse Wishart, given a restriction on a single diagonal element. Nobile (2000) provides a direct solution for generating such a draw in his Algorithm 3, which only requires simulating (and then properly manipulating) a series of chi square and normal random variates.

Step 4: Sampling $\boldsymbol{\psi}$

The conditional posterior distribution of the regression parameters $\boldsymbol{\psi}$ is multivariate normal:

$\boldsymbol{\psi} | \boldsymbol{\Xi}_{-\psi}, \boldsymbol{D}$

$$\sim N\left\{ \left[\boldsymbol{V}_\psi^{-1} + \sum_{i=1}^n \begin{pmatrix} \tilde{\boldsymbol{x}}_i & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \tilde{\boldsymbol{z}}_i & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \tilde{\boldsymbol{w}}_i \end{pmatrix}' \right.\right.$$
$$\left. \times \boldsymbol{\Sigma}_{c_i}^{-1} \begin{pmatrix} \tilde{\boldsymbol{x}}_i & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \tilde{\boldsymbol{z}}_i & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \tilde{\boldsymbol{w}}_i \end{pmatrix} \right]^{-1}$$
$$\times \left[\boldsymbol{V}_\psi^{-1}\boldsymbol{\psi}_0 + \sum_{i=1}^n \begin{pmatrix} \tilde{\boldsymbol{x}}_i & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \tilde{\boldsymbol{z}}_i & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \tilde{\boldsymbol{w}}_i \end{pmatrix}' \boldsymbol{\Sigma}_{c_i}^{-1} \begin{pmatrix} a_i^* \\ r_i^* \\ d_i^* \end{pmatrix} \right],$$
$$\left[\boldsymbol{V}_\psi^{-1} + \sum_{i=1}^n \begin{pmatrix} \tilde{\boldsymbol{x}}_i & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \tilde{\boldsymbol{z}}_i & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \tilde{\boldsymbol{w}}_i \end{pmatrix}' \right.$$
$$\left.\left. \times \boldsymbol{\Sigma}_{c_i}^{-1} \begin{pmatrix} \tilde{\boldsymbol{x}}_i & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \tilde{\boldsymbol{z}}_i & \boldsymbol{0}_{1\times k_{\tilde{w}}} \\ \boldsymbol{0}_{1\times k_{\tilde{x}}} & \boldsymbol{0}_{1\times k_{\tilde{z}}} & \tilde{\boldsymbol{w}}_i \end{pmatrix} \right]^{-1} \right\}.$$

Step 5: Sampling the latent $\boldsymbol{a}^*$

Eq. (15) reveals that each a_i^* can be sampled independently, $i = 1, 2, \ldots, n$, from the following density:

$$p(a_i^* | \boldsymbol{\Xi}_{-a_i^*}, \boldsymbol{D}) \propto \exp\left[-\frac{1}{2}(a_i^* - \mu_{a_i|u,v})^2 \sigma_{\epsilon\epsilon c_i|u,v}^{-1} \right]$$
$$\times [I(a_i = a_i^*) \times I(a_i^* < m_i) + I(a_i = m_i)$$
$$\times I(a_i^* \geq m_i)],$$

where

$$\mu_{a_i|u,v} = \tilde{\boldsymbol{x}}_i\tilde{\boldsymbol{\beta}} + \begin{pmatrix} \sigma_{\epsilon u c_i} & \sigma_{\epsilon v c_i} \end{pmatrix} \begin{pmatrix} \sigma_{uu c_i} & \sigma_{uv c_i} \\ \sigma_{vu c_i} & \sigma_{vv c_i} \end{pmatrix}^{-1}$$
$$\times \begin{pmatrix} r_i^* - \tilde{\boldsymbol{z}}_i\tilde{\boldsymbol{\gamma}} \\ d_i^* - \tilde{\boldsymbol{w}}_i\tilde{\boldsymbol{\theta}} \end{pmatrix}$$

$$\sigma_{\epsilon\epsilon c_i|u,v} = \sigma_{\epsilon\epsilon c_i} - \begin{pmatrix} \sigma_{\epsilon u c_i} & \sigma_{\epsilon v c_i} \end{pmatrix} \begin{pmatrix} \sigma_{uu c_i} & \sigma_{uv c_i} \\ \sigma_{vu c_i} & \sigma_{vv c_i} \end{pmatrix}^{-1}$$
$$\times \begin{pmatrix} \sigma_{u\epsilon c_i} \\ \sigma_{v\epsilon c_i} \end{pmatrix}.$$

Therefore, if $a_i < m_i$, $a_i^* = a_i$. If $a_i = m_i$, the latent log amount a_i^* is sampled from a univariate truncated normal distribution:

$a_i^* | \boldsymbol{\Xi}_{-a_i^*}, \boldsymbol{D} \sim \mathcal{TN}_{[m_i, \infty)}(\mu_{a_i|u,v}, \sigma_{\epsilon\epsilon c_i|u,v})$.

Step 6: Sampling the cutpoints $\boldsymbol{\delta}$

As has been well documented in the literature (e.g., Cowles, 1996 and Nandram and Chen, 1996), in samples of moderate size, simulation of the cutpoints $\boldsymbol{\delta}$ using standard Gibbs methods that condition on the latent response can mix very poorly. For this reason, blocking the cutpoints $\boldsymbol{\delta}$ and latent data $\boldsymbol{r}^*$ together in a single step can greatly improve the mixing performance of our posterior simulator. In what follows, therefore, we block these together using a marginal-conditional decomposition, first sampling the cutpoints $\boldsymbol{\delta}$.

Marginalized over $\boldsymbol{r}^*$, the conditional posterior distribution of $\boldsymbol{\delta}$ is proportional to

$$p(\boldsymbol{\delta} | \boldsymbol{\Xi}_{-\delta, r^*}, \boldsymbol{D}) \propto I(\delta_1 = -\infty)I(\delta_2 = 0)I(0 < \delta_3 < \delta_4 < \delta_5 < 1)$$
$$\times I(\delta_6 = 1)I(\delta_7 = \infty)$$
$$\times \prod_{i=1}^n [\Phi_{r_i}(\delta_{r_i+2}) - \Phi_{r_i}(\delta_{r_i+1})],$$

where $\Phi_{r_i}(\cdot)$ denotes the cumulative distribution function of a normal density with mean $\mu_{r_i|\epsilon,v}$ and variance $\sigma_{uu c_i|\epsilon,v}$ where

$$\mu_{r_i|\epsilon,v} = \tilde{\boldsymbol{z}}_i\tilde{\boldsymbol{\gamma}} + \begin{pmatrix} \sigma_{u\epsilon c_i} & \sigma_{uv c_i} \end{pmatrix}$$
$$\times \begin{pmatrix} \sigma_{\epsilon\epsilon c_i} & \sigma_{\epsilon v c_i} \\ \sigma_{v\epsilon c_i} & \sigma_{vv c_i} \end{pmatrix}^{-1} \begin{pmatrix} a_i^* - \tilde{\boldsymbol{x}}_i\tilde{\boldsymbol{\beta}} \\ d_i^* - \tilde{\boldsymbol{w}}_i\tilde{\boldsymbol{\theta}} \end{pmatrix}$$
$$\sigma_{uu c_i|\epsilon,v} = \sigma_{uu c_i} - \begin{pmatrix} \sigma_{u\epsilon c_i} & \sigma_{uv c_i} \end{pmatrix}$$
$$\times \begin{pmatrix} \sigma_{\epsilon\epsilon c_i} & \sigma_{\epsilon v c_i} \\ \sigma_{v\epsilon c_i} & \sigma_{vv c_i} \end{pmatrix}^{-1} \begin{pmatrix} \sigma_{\epsilon u c_i} \\ \sigma_{vu c_i} \end{pmatrix}.$$

As this does not correspond to any standard density, we employ an M–H step to sample values from this conditional posterior. Specifically, following Nandram and Chen (1996) and Li and Tobias (2008), we generate candidate values for the cutpoint vector by first sampling their differences $\boldsymbol{q}$ from a Dirichlet proposal density:

$\boldsymbol{q}^* = (q_1^*, q_2^*, q_3^*, q_4^*)$
$\quad \sim \text{Dirichlet}(\alpha_1 l_1 + 1, \alpha_2 l_2 + 1, \alpha_3 l_3 + 1, \alpha_4 l_4 + 1)$

where $*$ indicates a candidate draw, $q_j^* \equiv \delta_{j+2}^* - \delta_{j+1}^*$, α_j are tuning parameters and $l_j = \sum_{i=1}^n I(r_i = j)$ for $j = 1, 2, 3, 4$. We accept

the candidate δ^* with probability

$$\min\left\{1, \left[\prod_{i=1}^{n}\frac{\Phi_{r_i}(\delta^*_{r_i+2}) - \Phi_{r_i}(\delta^*_{r_i+1})}{\Phi_{r_i}(\delta^\dagger_{r_i+2}) - \Phi_{r_i}(\delta^\dagger_{r_i+1})}\right]\left[\prod_{j=1}^{4}\left(\frac{q_j^\dagger}{q_j^*}\right)^{\alpha_j l_j}\right]\right\}$$

where † indicates the last accepted draw from the M–H step. In our application, we set $\alpha_j = 0.1 \ \forall j$, which leads to acceptance rates approximately equal to 10%. Inefficiency factors for our cutpoint parameters ranged from 27 to 29, with effective sample sizes approximately equal to 3000.

Step 7: Sampling the latent renewals variable $\boldsymbol{r}^*$

Eq. (15) reveals that the latent renewals r_i^* can be sampled independently, $i = 1, 2, \ldots, n$, from the following density, given the values of the cutpoints $\boldsymbol{\delta}$ just sampled:

$$r_i^* | \boldsymbol{\Xi}_{-r_i^*}, \boldsymbol{D} \sim \mathcal{TN}_{(\delta_{r_i+1}, \delta_{r_i+2})}(\mu_{r_i|\epsilon,v}, \sigma_{uuc_i|\epsilon,v}), \quad i = 1, 2, \ldots, n.$$

Step 8: Sampling the latent default $\boldsymbol{d}^*$

Finally, each latent default value d_i is sampled independently from the following univariate truncated normal distribution:

$$d_i^* | \boldsymbol{\Xi}_{-d_i^*}, \boldsymbol{D} \sim \begin{cases} \mathcal{TN}_{(-\infty,0]}(\mu_{d_i|\epsilon,u}, \sigma_{vvc_i|\epsilon,u}) & \text{if } d_i = 0 \\ \mathcal{TN}_{(0,\infty)}(\mu_{d_i|\epsilon,u}, \sigma_{vvc_i|\epsilon,u}) & \text{if } d_i = 1 \end{cases}$$

where

$$\begin{aligned}
\mu_{d_i|\epsilon,u} &= \tilde{\boldsymbol{w}}_i\tilde{\boldsymbol{\theta}} + \begin{pmatrix}\sigma_{v\epsilon c_i} & \sigma_{vuc_i}\end{pmatrix}\begin{pmatrix}\sigma_{\epsilon\epsilon c_i} & \sigma_{\epsilon uc_i} \\ \sigma_{u\epsilon c_i} & \sigma_{uuc_i}\end{pmatrix}^{-1} \\
&\quad \times \begin{pmatrix}a_i^* - \tilde{\boldsymbol{x}}_i\tilde{\boldsymbol{\beta}} \\ r_i^* - \tilde{\boldsymbol{z}}_i\tilde{\boldsymbol{\gamma}}\end{pmatrix} \\
\sigma_{vvc_i|\epsilon,u} &= \sigma_{vvc_i} - \begin{pmatrix}\sigma_{v\epsilon c_i} & \sigma_{vuc_i}\end{pmatrix} \\
&\quad \times \begin{pmatrix}\sigma_{\epsilon\epsilon c_i} & \sigma_{\epsilon uc_i} \\ \sigma_{u\epsilon c_i} & \sigma_{uuc_i}\end{pmatrix}^{-1}\begin{pmatrix}\sigma_{\epsilon vc_i} \\ \sigma_{uvc_i}\end{pmatrix}.
\end{aligned}$$

References

Albert, J.H., Chib, S., 1993. Bayesian analysis of binary and polychotomous response data. Journal of the American Statistical Association 88 (422), 669–679.

Chib, S., Jacobi, L., 2007. Modeling and calculating the effect of treatment at baseline from panel outcomes. Journal of Econometrics 140, 781–801.

Cowles, M., 1996. Accelerating Monte Carlo Markov chain convergence for cumulative-link generalized linear models. Statistics and Computing 101–111.

Deb, P., Munkin, M.K., Trivedi, P.K., 2006. Bayesian analysis of the two-part model with endogeneity: application to health care expenditure. Journal of Applied Econometrics 21, 1081–1099.

Elliehausen, G., Lawrence, E., 2001. Payday advance credit in America: an analysis of customer demand. Georgetown University Credit Research Center Monograph No. 35.

Flannery, M.J., Samolyk, K., 2005. Payday lending: do the costs justify the price? FDIC Center for Financial Research Working Paper No. 2005/09.

Frühwirth-Schnatter, S., 1994. Data augmentation and dynamic linear models. Journal of Time Series Analysis 15, 183–202.

Frühwirth-Schnatter, S., 2001. Markov chain Monte Carlo estimation of classical and dynamic switching and mixture models. Journal of the American Statistical Association 96, 194–209.

Frühwirth-Schnatter, S., 2004. Estimating marginal likelihoods for mixture and Markov switching models using bridge sampling techniques. Econometrics Journal 7, 143–167.

Geweke, J., 2007. Interpretation and inference in mixture models: simple MCMC works. Computational Statistics and Data Analysis 3529–3550.

Geweke, J., Gowrisankaran, G., Town, R.J., 2003. Bayesian inference for hospital quality in a selection model. Econometrica 71 (4), 1215–1283.

Jeliazkov, I., Graves, J., Kutzbach, M., 2008. Fitting and comparison of models for multivariate ordinal outcomes. Advances in Econometrics: Bayesian Econometrics 23, 115–156.

Kleibergen, F., van Dijk, H.K., 1998. Bayesian simultaneous equations analysis using reduced rank structures. Econometric Theory 14, 701–743.

Kloek, T., van Dijk, H.K., 1978. Bayesian estimates of equation system parameters: an application of integration by Monte Carlo. Econometrica 46 (1), 1–19.

Lancaster, A., 2004. An Introduction to Modern Bayesian Econometrics. Blackwell Publishing.

Li, K., 1998. Bayesian inference in a simultaneous equation model with limited dependent variables. Journal of Econometrics 85 (2), 387–400.

Li, M., Tobias, J., 2008. Bayesian analysis of treatment effects in an ordered potential outcomes model. In: Advances in Econometrics, Volume 21:Estimating and Evaluating Treatment Effects in Econometrics.

McCulloch, R., Polson, N., Rossi, P., 2000. A Bayesian of the multinomial probit model with fully identified parameters. Journal of Econometrics 173–193.

Melzer, B.T., 2011. The real costs of credit access: evidence from the payday lending market. Quarterly Journal of Economics 126 (1), 517–555.

Morse, A., 2011. Payday lenders: heroes or villains?. Journal of Financial Economics 102 (1), 28–44.

Munkin, M.K., Trivedi, P.K., 2003. Bayesian analysis of a self-selection model with multiple outcomes using simulation-based estimation: an application to the demand for health care. Journal of Econometrics 114 (2), 197–220.

Munkin, M.K., Trivedi, P.K., 2008. Bayesian analysis of the ordered probit model with endogenous selection. Journal of Econometrics 143 (2), 334–348.

Nandram, B., Chen, M.-H., 1996. Reparameterizing the generalized linear model to accelerate Gibbs sampler convergence. Journal of Statistical Computation and Simulation 129–144.

Nobile, A., 2000. Comment: Bayesian multinomial probit models with a normalization constraint. Journal of Econometrics 99 (2), 335–345.

Skiba, P.M., Tobacman, J., 2007. The profitability of payday loans. Working Paper.

Skiba, P.M., Tobacman, J., 2008. Payday loans, uncertainty and discounting: explaining patterns of borrowing, repayment, and default. Vanderbilt Law and Economics Research Paper No. 08-33.

Skiba, P.M., Tobacman, J., 2009. Do payday loans cause bankruptcy? SSRN Working Paper 1266215.

Tanner, M.A., Wong, W.H., 1987. The calculation of posterior distributions by data augmentation. Journal of the American Statistical Association 82 (398), 528–540.

Zellner, A., L, Bauwens, van Dijk, H.K., 1988. Bayesian specification analysis and estimation of simultaneous equations models using Monte Carlo methods. Journal of Econometrics 38, 39–72.