Response to Begenau and Stafford (2023)

May 16, 2024

1 Summary

Begenau and Stafford (2023; henceforth, “BS”) attempt to replicate Drechsler, Schnabl, and Savov (QJE, 2017; henceforth, “DSS”). They acknowledge that they can replicate the main results from the paper.

However, BS claim that some DSS results are not fully robust when expanding the estimation sample. Specifically, BS argue that some of the DSS regressions should also include data on so-called follower branches. Follower branches are branches that do not set their own deposit rates according to the data provider RateWatch. Instead, follower branches adopt the deposit rate from another branch.

BS make two main claims with respect to follower branches:

- BS Claim 1: The deposit flow results hold for follower branches.
  - Specifically, the coefficient of interest in DSS Table 3, Column 1, is statistically significant and negative when estimated for follower branches only.
  - BS claim this shows that a ”placebo regression” fails, i.e., BS argue that deposit flows respond to interest rates even for follower branches, whose rates are not necessarily set locally.
- BS Claim 2: The deposit rate results disappear when including follower branches.
  - Specifically, the coefficient of interest in DSS Table 2, Column 1, is statistically insignificant if the regression includes follower branches.
  - BS claim this shows that one of the DSS results is not robust when considering a broader sample.

In this report, we examine both BS claims and find that both are incorrect. Specifically, this report shows the following results:

- Result 1: BS Claim 1 is incorrect because BS do not use the correct data. BS backfill their data. The results do not hold when estimating the regressions without incorrectly backfilling the data.
- Result 2: BS Claim 2 is incorrect due to a coding error made by BS.
- Result 3: BS make a conceptual error in testing BS Claims 1 and 2. Specifically, BS ignore that uniform pricing means that, within rate-setting families, one must consider the local market conditions of *all* branches, including follower branches. We show that the DSS results are robust to including follower branches when properly accounting for them within a rate-setting family.

BS also make several other secondary claims regarding (i) the relationship between market power and uniform pricing, (ii) the distinction between market concentration and market power, (iii) whether the deposits channel applies to large banks, (iv) the scope of the DSS research design, (v) the lack of an alternative explanation, and (vi) various other issues. We demonstrate that these claims are also incorrect.

2 Response

2.1 BS replication using follower branches

2.1.1 BS replication fails because of coding errors

BS confirm that they can replicate the DSS results. Hence, there is no disagreement over the original DSS results.

However, BS claim that DSS do not properly account for follower branches. They argue that the DSS results are not robust (see BS Claims 1 and 2) when including follower branches. We conduct a replication of the BS replication and find that the BS claims are due to incorrect backfilling of data and coding errors. To give BS their best shot, we use the BS code for our replication.

BS Claim 1 is based on incorrect backfilling of follower branches

BS claim that they find a statistically significant result for follower branches in the deposit flows regression. They consider this a "placebo regression" that should deliver an insignificant result.

At it turns out, BS make a fundamental error when constructing their dataset of follower branches.

The issue is straightforward. Ratewatch provides a file that matches ratesetters to follower branches. Importantly, this file specifies the date on which a follower branch starts following a specific ratesetter. As it turns out, the early years of the Ratewatch data only provide very limited coverage of follower branches.

BS artificially increase the number of followers by *backfilling* this ratesetter-follower link. For example, suppose a branch is designated a follower branch starting in January 2013 and does not appear in the dataset prior to that. BS simply assume that the follower was already a follower in prior years and backfill the data to Jan 2001.

It is obvious that backfilling can lead to biased results. If some branches set their own deposit rates during the period in which BS incorrectly classifies them as followers, then this will generate the "placebo" effect. This is plausible given the large increase in the follower branches during the analysis period (see Figure 1 below). It is also consistent with analysis using Ratewatch data in Granja and Paixao (2019), who argue that *"uniform pricing practices have become more prevalent over time"*. In other words, backfilling can generate the "placebo" effect.

Here is the relevant part of the BS *"ReadME"* documentation that confirms that BS are backfilling their data:

```
makeSODRWcombinePanel.do (in Stata) makes a panel of SOD data that has been
merged with RatWatch data, backfilling links. We use this for
DepFlow_FollowerPlacebo.do
        o Input is ratewatch_panel_MM25K.dta; sodbranch_sample.dta
        o Saves "`dir_output'\branch_ratewatch_sample.dta", in
          processed data folder
```

The BS code implements the backfilling in lines 57-70 of the file *"makeSODRWcombinePanel.do"*.

Figure 1: **Ratewatch follower branches coverage**

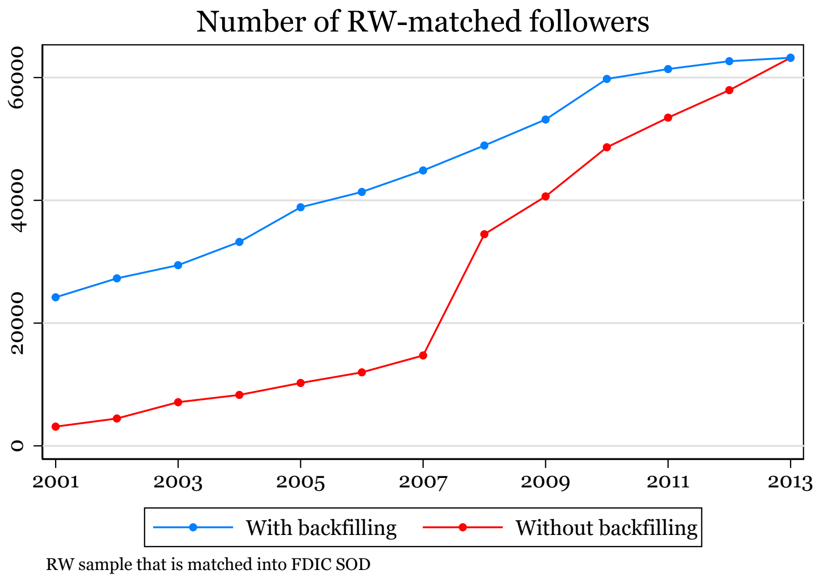

Figure 1 shows the total number of follower branches per year with and without backfilling. We note that the backfilled data contain 543,188 observations. However, the correct data only contain 330,829 observations. Hence, BS almost double the dataset by backfilling observations.

We note that in the correct (non-backfilled) data, the vast majority of follower observations occur after the year 2008. Importantly, there is no variation in the Fed funds rate after 2008. Hence, in the correct data the estimates are not particularly sensitive to follower branches because they are less relevant for identification. Yet, the estimates are quite sensitive to follower branches in the backfilled data because the backfilled observations mostly cover the years before 2008.

Moreover, we emphasize that BS estimate the flow regression for the period from 2001 to 2013 (see Table 7 Panel B). This is a significant departure from DSS who estimate the regression for the years 1994 to 2013. This is important because the deposit flow dataset is annual and the coefficient of interest requires variation in the Fed funds target rate. Yet, the Fed funds rate only varied from 1994 to 2008 and was effectively unchanged thereafter. Hence, BS effectively cut the number of years with variation in the Fed Funds rate in half, thereby relying on a selected part of the dataset.

We can now compare the results with and without backfilling. We use the replication data and code posted by DSS on the QJE website as of 2017. In addition, we use the ratesetter-follower link provided by Ratewatch. This approach uses the original DSS data and is straightforward to replicate.

Table 1: **BS Table 7 Panel B replication, followers only**

	Δlog Deposits, followers only					
	2001-2013			2001-2009		
	BS	Replication		BS	Replication	
Backfilled (YES/NO):	YES	YES	NO	YES	YES	NO
	(1)	(2)	(3)	(4)	(5)	(6)
ΔHHI x dFFR	−0.70**	−0.57	0.06	−0.93**	−0.91**	−0.43
	(0.33)	(0.41)	(0.50)	(0.37)	(0.42)	(0.54)
Bank × year FE	Yes	Yes	Yes	Yes	Yes	Yes
State × year FE	Yes	Yes	Yes	Yes	Yes	Yes
Branch FE	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
County × ZLB FE	Yes	Yes	Yes	No	No	No
Obs.	607,738	543,188	330,829	351,236	306,842	118,589
R^2	0.399	0.318	0.345	0.468	0.392	0.441

We reproduce the BS result using the BS code in Column (1) of Table 2 (replicates BS Table 7, Panel B, Column (2)). In Column (2) we use the DSS data and construct a dataset that is incorrectly backfilled, following the BS backfilling procedure. We find a coefficient of −0.57, which is similar to the BS coefficient reported in Column (1). Contrary to BS, our coefficient is *not* statistically significant. The difference in coefficients between Columns (1) and (2) is however due to various BS changes relative to the original DSS sample. As discussed in more detail below, BS make a number of changes to the original data that affect the estimation (different branch identifier, including institutions other than commercial banks, different definition of branch fixed effects, etc.).

Column (3) presents the results without the incorrect backfilling. Now the coefficient is 0.06 and statistically insignificant. Hence, the "placebo" effect disappears once we correct for the erroneous backfilling. As discussed above, the backfilling may incorrectly designate branches that set their own rate to be follower branches. This will generate a "placebo" effect when incorrectly backfilling the data.

Columns (4) to (6) present results for the period from 2001 to 2009. BS use this time period even though it is different from the one considered in DSS. Again, the results disappear once we correct for the erroneous backfilling. Column (4) presents the replication of BS using their code and backfilled data (replicates BS Table 7, Panel B, Column (4)). Column (5) uses the DSS data and backfilling, which replicates the backfilled results. Column (6) presents the results without backfilling. Once more, the "placebo" result disappears once we do not backfill the data.

BS never acknowledge or discuss backfilling in their paper and provide no justification for doing so. The backfilling can only be inferred from looking at their code. In fact, the first version of this paper, Begenau and Stafford (2022), also reported regressions with backfilled data. This is apparent when carefully looking at Begenau and Stafford (2022). Begenau and Stafford (2022) state that they only had Ratewatch data starting in 2001. Moreover, Begenau and Stafford (2022) claim to run follower regressions for a matched Ratewatch-FDIC sample. Begenau and Stafford (2022) then report flow regression covering the years 1994 to 2008 (see Begenau and Stafford (2022), Table 4, Panel B). Yet, it is impossible to have a matched FDIC-Ratewatch data for the years 1994 to 2000 given that Ratewatch starts in 2001. Hence, it follows that Begenau and Stafford (2022) also backfilled observations.

To be clear, there is no reason to believe that the Ratewatch data can or should be backfilled. Ratewatch specifically states that a ratesetter-follower link only starts on the date provided by Ratewatch, inconsistent with backfilling. We contacted Ratewatch directly and verified this issue when we first worked on DSS. WE also called them again after reading the BS paper to verify the same issue again. Ratewatch explicitly confirmed that they have no data on the link between ratesetter and followers beyond what is included in their matching file. We were told that backfilling with the BS procedure is equivalent to making up data.

Therefore, we reject BS claim 1.

BS Claim 2 is based on incorrect coding of bank ownership

To give BS their best shot, we implement their code and data to replicate the BS results.

BS report their main results on deposit rates using the DSS sample in Table A4. This is the replication that most closely matches the DSS specification and sample. BS report these results in the Appendix. BS also offer a second "replication" that is discussed in the main text (Table 6) and that differs from DSS in several important ways. We discuss BS Table 6 below.

We start by replicating their main result using their code and data. We replicate the relevant columns (Columns 1 and 5) from Table A4. The replication exactly matches the results reported in BS. Column (1) of Table 2 reports the result for the sample that includes ratesetting branches and follower branches (Column 1 in Table A4). We replicate the statistically significant coefficient of 0.05 reported in BS. Column (2) replicates the BS result for ratesetters only (Column 5 in Table A4). We replicate the statistically significant coefficient of 0.12.

Table 2: **BS Table A4 replication**

	ΔSavings rate spread					
	BS results		Replication			
	All	Ratesetters	All	Ratesetters	All	Ratesetters
	(1)	(2)	(3)	(4)	(5)	(6)
ΔFF × branch HHI	0.05***	0.12***	0.02**	0.06**	0.02**	0.07**
	(0.01)	(0.02)	(0.01)	(0.03)	(0.01)	(0.03)
Bank × quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
State × quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Branch FE	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
County × ZLB FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	1,481,416	251,763	1,182,139	113,663	1,155,566	87,892
R^2	0.872	0.698	0.931	0.788	0.946	0.861

Importantly, we note that BS's own replication rejects BS claim 2. Instead, consistent with DSS, Column (1) shows that branches located in high-concentration counties increase deposit spreads more (i.e., increase deposit rates less) when the Fed funds rate increases relative to branches in low-concentration counties, even when including follower branches in the sample. Hence, BS show that the DSS result is robust to including follower branches.

BS downplay their own finding. Although they acknowledge that this result supports DSS, they try to dismiss it. Their reasoning goes as follows (see their discussion on page 55). They argue that Ratewatch only provides a single cross-section of bank-branch ownership and construct bank-time fixed effects from this cross-section. BS point out that using a single cross-section leads to mismeasurement of bank-time fixed effects because of incorrect bank-branch ownership (e.g., due to mergers). They assert that using the correct bank branch ownership would cause the DSS result to "disappear." Curiously, BS never show that the result indeed "disappears" after correcting bank-branch ownership.

We reject the claim that DSS used incorrect bank ownership. DSS were aware that Ratewatch only provided a single cross-section of bank-branch ownership. This was one of the reasons DSS matched the Ratewatch data to FDIC data. DSS used FDIC data to correctly assign bank branch ownership over time. It is therefore straightforward to rerun the BS regression using bank-time fixed effects that are based on correct bank ownership data.

Columns (3) and (4) of Table 2 present the results. The only change relative to Columns (1) and (2) is that we construct bank-time fixed effects correctly.

We do *not* find that the DSS results "disappear." Column (3) shows a statistically significant coefficient of 0.02 for the sample including ratesetters and followers. Column (4) finds a statistically significant coefficient of 0.06 for the sample of ratsetters only. This shows that the DSS result is robust once we correct for the BS coding error.

<u>Therefore, we reject BS Claim 2.</u>

We make four observations on these results:

1. The coefficient of interest tends to decrease when including follower branches. The relative decline is the same in the original BS specification (compare Columns 1 and 2) and our corrected specification (compare Columns 3 and 4). This is mechanical given how the specification is set up. It comes from the fact that the left-hand side variable (deposit spreads) does not vary within a rate-setting family, while the right-hand side does (county HHI). This is the reason why simply adding follower branches to the DSS specification is not the correct way to evaluate the impact of follower branches. We discuss how to properly address this issue in the second part of our report.

2. BS sometimes claim that the data before 2001 is unreliable. We disagree with this characterization. In any case, we can estimate the main specifications after dropping the years 1997 to 2000. We present the results in Column (5) and (6) of Table 2. The coefficients of interest are almost unchanged relative to Columns (3) and (4), respectively.

3. We note that the sample size decreases when comparing Columns (1) and (2) with Columns (3) and (4), respectively. The reason is that some Ratewatch branches cannot be matched to FDIC data. Note that after we exclude non-matched branches, we find a lower coefficients for specifications *with and without* follower branches. Hence, the coefficient of interest in part declines because we work with a smaller, selected dataset – not because of changes in the coding of bank ownership.

4. BS also discuss a rate regression in which the main coefficient of interest is -0.01 and statistically insignificant (see Table 6, Column (1)). We note that this regression uses a sample, a time period, and variable definitions significantly different from DSS. BS do not document or discuss these differences and they could only be inferred by analyzing the BS code. As far as we can tell, the only common denominator for these differences appears to be that they tend to lower the coefficient of interest.

   Here is an overview of differences between BS and DSS that we have identified from the BS code. There may be other differences to the original DSS results that BS have not documented and that we have not identified. Some of these differences also apply to the regression in Table A4.

   (a) Inconsistent branch identifier: Ratewatch provides a branch-level file to link Ratewatch data to FDIC data. This is the file used in DSS. Some Ratewatch branches cannot be matched to FDIC data based on the Ratewatch branch identifier. BS use a fuzzy matching algorithm to match Ratewatch and FDIC branches using location data. The algorithm matches branches that are not otherwise matched and sometimes also overrides the matching file provided by Ratewatch. There is only a brief description of the issue in the paper (see page 6).

After applying the fuzzy matching procedure, BS end up with two different branch identifiers, referred to as "soduninumbr" and "uninumbrlookup", respectively, as well as two bank identifiers, referred to as "sodcert" and "rwcert". There is a significant number of branches for which the resulting branch and bank identifier differ (i.e., "soduninumbr" is different from "uninumbrlookup"). Here are examples of branch observations with two different branch identifiers from the BS data *"ratewatch_panel_MM25K.dta"*:

rwcert	uninumbrlookup	sodcert	soduninumbr	date_q
24897	32123	32633	282932	2002q1
9087	359069	27714	274924	2002q1
19629	422460	3180	3770	2002q1
15277	2239	3511	188848	2002q1

Subsequently, BS use one set of identifiers for constructing and cleaning the data in some of their code (see bank and branch identifiers in lines 71-75 of *"makePanelForReg.do"*). BS also use the other set of identifiers for sample restrictions and estimation in other parts of their code (see bank and branch identifiers in lines 16-19 in *"DSStab2_replication_paper.do"*).

Obviously, this approach of using different identifiers in different parts of the code is not consistent. Moreover, it is unclear why this matching procedure would be an improvement over using the original Ratewatch matching file.

(b) Data processing errors: BS have issues with loading some of their 2015 version Ratewatch data due to several coding errors:

- As a result of incorrectly dropping duplicate observations in a quarter, BS keep some mid-quarter (instead of end-of-quarter) rates for branches.
- BS have missing data in 2012 due to some rates being loaded as strings, as shown in this screenshot:

2011q1	58,287	2.32	51.02
2011q2	60,141	2.40	53.41
2011q3	61,275	2.44	55.85
2011q4	62,279	2.48	58.34
2012q1	63,514	2.53	60.87
2012q2	64,027	2.55	63.42
2012q3	$\boxed{419}$	0.02	63.44
2012q4	66,612	2.65	66.09
2013q1	69,168	2.76	68.85
2013q2	71,450	2.85	71.69
2013q3	72,860	2.90	74.60
2013q4	74,230	2.96	77.55

- The Ratewatch ratesetter-follower link starts in 1999. BS's code changes all start years from 1999 to 1997 (variable name *year*) without changing their date variable (variable name *date_des*). This causes the constructed *month* variable for these changed observations to be over 200, and their following *date_d* variable to be missing. As a result of this mistake, these observations are dropped. The screenshot shows the re-coding of the year.

```
replace year =1997 if year ==1999 // otherwise all links established after files exist

gen month = floor(date_des) - year*100
drop date_des temp
gen date_d = mdy(month, day,year)
format date_d  %td
```

(c) Different sample: DSS include all commercial banks in their analysis. BS expand the sample to include institutions other than commercial banks. BS provide no justification for this change relative to the original paper and do not discuss it in the paper.

(d) Construction of branch identifier: DSS use the branch identifier provided by the FDIC. BS construct a new identifier based on the interaction of bank identifier and branch identifier. This significantly increases the number of identifiers. As a result, BS have a significantly larger number of fixed effects, which leads BS to drop additional observations and changes the main DSS specification. Moreover, as discussed above, BS have two separate sets of identifiers. BS provide no justification for this change to the original paper and do not discuss it in the paper.

(e) Different definition of HHI: DSS provide county-level HHI as part of the QJE replication package. BS construct alternative versions of HHI, including a time-varying version and a version including institutions other than commercial banks. Both HHI variables are used in BS. BS provide no justification for this change relative to the original paper and do not discuss it in the paper.

(f) Different time period: BS cover a different time period. DSS covers Jan 1997 to Dec 2013. Table 6 only covers Apr 2001 to Dec 2013.

(g) Additional coding errors: BS themselves identified an additional coding error after submitting their paper. This error is different from the ones discussed above.

In short, BS make several changes in Table 6 relative to DSS that are neither justified nor discussed. They also make several coding errors. We believe a proper replication should focus on the issue at hand (how to account for follower branches) instead of changing the original specification in different ways. BS's Table 6 therefore fails to replicate DSS.

2.1.2 BS replication fails because of a conceptual error

Aside from the discussion above, there is also a conceptual issue with how BS account for follower branches. Even if BS Claims 1 and 2 were correct, the empirical tests for these claims would be incorrect. The BS analysis of follower branches are not a proper robustness test of DSS.

The conceptual issue arises because of the rationale behind setting up ratesetting families. Ratesetting families tend to be geographically clustered. However, there is significant variation across families in terms of their coverage. Some families are concentrated in a few counties, while others cover much larger areas.

For the sake of the argument, let us assume that one wants to test DSS including follower branches. If one includes follower branches, one needs to account for how banks set deposit rates in a ratesetting family. Economic intuition suggests that a profit-maximizing bank sets an optimal deposit rate for a family considering market conditions across all branches within the family.

It follows that market conditions of a family cannot simply be measured by the market conditions for a single branch. Yet, this is exactly what BS do. BS assign the county-level HHI to each branch based on where the branch is located, while ignoring its relationship to other branches within the same family.

As a result, adding follower branches to the DSS specification mechanically decreases the coefficient of interest. Intuitively, BS add observations for which the right-hand side variable varies, while the left-hand side variable does not, thereby mechanically lowering the coefficient.

There is a simple and straightforward approach to address this issue. To include follower branches, one must recognize that profit-maximizing banks need to set their deposit rates with respect to market concentration (i.e., HHI) across all markets in which a family is operating. We therefore compute the average HHI for a ratesetting family across all branch locations within the family. Intuitively, this approach incorporates the average market concentration (HHI) in all of the family's locations. By revealed preference, this is the relevent market as chosen by the ratesetting family.

To implement this approach, we use the Ratewatch ratesetter-follower link file to generate the sample that includes ratesetter and follower branches. We compute an equal-weighted HHI across all branches of a family. In addition, we compute a value-weighted HHI across branches of a family using lagged branch deposits as weights. We then estimate the same regression and same specification as in DSS. The difference to the original DSS specification is that an observation now represents a ratesetting family instead of a ratesetting branch. This approach accounts for differences across families in terms of the relevant market.

Table 3: **Rates family regressions**

	ΔSavings rate spread					
	Equal-weighted		Value-weighted		WLS	
	(1)	(2)	(3)	(4)	(5)	(6)
ΔFF × avg fam-HHI	0.10***	0.15***	0.10***	0.15***	0.10**	0.15***
	(0.03)	(0.02)	(0.03)	(0.02)	(0.05)	(0.02)
Bank × time FE	Yes	No	Yes	No	Yes	No
State × time FE	Yes	Yes	Yes	Yes	Yes	Yes
Branch FE	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
County × ZLB FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	117,926	414,893	117,926	414,893	117,926	414,893
R^2	0.821	0.664	0.821	0.664	0.877	0.655

Table 3 shows the results for the deposit rate regressions at the family-level. Columns (1) and (2) show the result using equal-weighted family HHI. Column (1) finds a statistically significant coefficient of 0.10 in the specification with the full set of fixed effects. Column (2) finds a statistically significant coefficient 0.15 in the specification without bank-time fixed effects. Columns (3) and (4) show the result using the value-weighted family HHI. We find identical coefficients of 0.10 and 0.15, respectively. We can also consider the impact of giving a larger weight to larger families. We weigh families by the total number of branches within the family relative to the total number of branches in a given quarter. Hence, the weights are proportional to family size. Columns (5) and (6) present the results and find the same coefficients. Moreover, these estimates are similar to the original DSS estimates, which vary between 0.10 and 0.20 (see DSS Table 2, Panel A).

Next, we examine the effect on deposit flows. We use the same specification and data as in DSS Table 3, Column (1). We compute deposit growth at the family level. Similar to the rates regression, we replace all individual branch-level observations belonging to the same ratesetting family with a single family-level observation. We use the same family-level HHIs (equal-weighted and value-weighted) as for the rate regressions.

Table 4: **Flows family regressions**

	Δlog Deposits					
	Equal-weighted		Value-weighted		WLS	
	(1)	(2)	(3)	(4)	(5)	(6)
ΔFF × avg fam-HHI	−0.93**	−2.48***	−0.97**	−2.49***	−0.87**	−2.36***
	(0.40)	(0.29)	(0.40)	(0.29)	(0.39)	(0.28)
Bank × year FE	Yes	No	Yes	No	Yes	No
State × year FE	Yes	Yes	Yes	Yes	Yes	Yes
Family FE	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
County × ZLB FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	895,661	957,227	895,661	957,227	895,661	957,227
R^2	0.329	0.220	0.329	0.221	0.327	0.218

Table 4 presents the results. Like in the original DSS sample, these regressions also include branches that are not covered by Ratewatch. Columns (1) and (2) use equal-weighted family HHI and find statistically significant coefficients of −0.93 and −2.48 for the specification with the full set of fixed effects and without bank-time fixed effects respectively. Columns (3) and (4) display the result using value-weighted family HHI, finding similar coefficients of −0.97 and −2.49 respectively. Finally, as shown in Columns (5) and (6), when we weigh families proportional to their size, the results remain similar with coefficients −0.87 and −2.36 respectively. These estimates are similar to the original DSS estimates, which vary between −0.66 and −1.82 (see DSS Table 3).

To summarize, we find that DSS results are robust to including follower branches when taking into account family structure.

2.1.3 Should BS add follower branches to the rate regressions?

The results above support the DSS results when accounting for follower branches. That said, let us briefly discuss why DSS did not originally include follower branches in the rates regressions.

As mentioned in DSS, one reason was that follower branches are duplicating existing observations. A concern was that by duplicating branches, we are mechanically lowering the coefficient of interest.

Another reason why DSS did not include follower branches was the rapidly changing Ratewatch coverage of follower branches. Figure 2 plots the number of followers branches at the time when they first enter the Ratewatch dataset. Note that there is very little entry during the first few years. The vast majority of followers enter the dataset in the years 2007 and 2008. This analysis raised two issues at the time of writing.

Figure 2: **Ratewatch follower branches coverage**

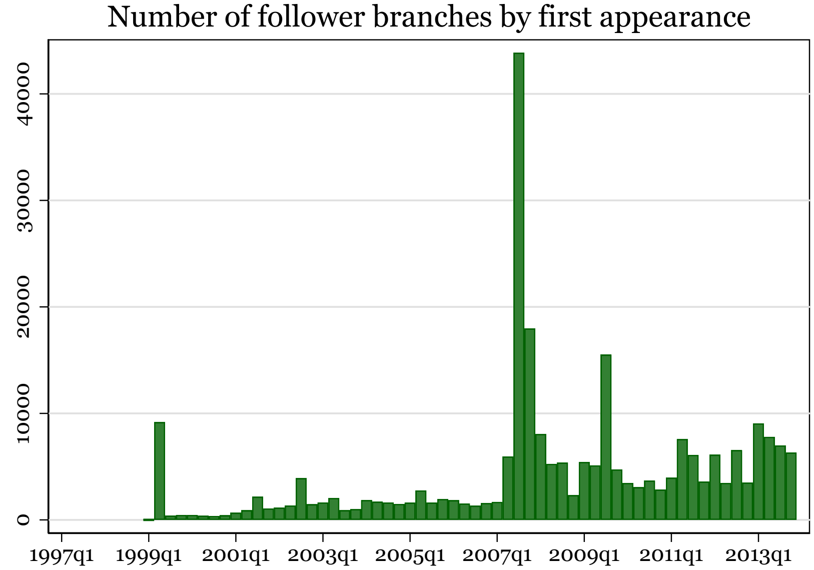

First, there is potential selection in terms of the coverage of follower branches. There was a concern that Ratewatch personnel coded branches close to a ratesetting branch as a follower if they tended to pay the same deposit rate. However, it was not necessarily confirmed that they were part of the same family. It was suggested to us that some branches may have been coded as followers even if there was no direct confirmation from the branch. As discussed above, this can generate the "placebo" similar to backfilling data.

Second, the Fed funds rate was unchanged after 2008. Given that follower coverage was quite limited before 2007 (see Figures 1 and 2), there was little right-hand side variation for most of the follower data. Indeed, as discussed above, the results are insignificant if the data is not backfilled.

BS do not discuss the dramatic change in coverage over the analysis period. Instead, they assume it away by backfilling the data. This is obviously incorrect.

So that there is no misunderstanding, we note that DSS always included all branches from the FDIC data in deposit flow regressions, independent of whether the branch was covered by Ratewatch or not. The reason is that there is no selection in the FDIC data because the FDIC data cover the entire universe and there is no concern about data reliability. This was preferred because the FDIC data are at a low frequency (annual instead of quarter) and we deal with quantities which tend to be less well-measured than prices. As discussed at length in the original review process, using the entire FDIC data back to 1994 therefore helps with statistical power.

2.2 Additional BS claims

BS make a number of other claims in their paper. Next, we provide a brief discussion of the main ones.

2.2.1 Lack of an alternative explanation

It takes a theory to beat a theory. BS explicitly acknowledge that they fail to provide an alternative theory that explains the findings in DSS. They respond to the following question: *"What are the omitted variables generating the deposit channel?"* with *"This is an important question that we approached in several ways. In the end, we do not have a fully satisfying answer to this question."* In short, they do not have an answer.

2.2.2 Market power and uniform pricing

BS mention market power 61 times in their draft. BS argue that their empirical results are important because they think that uniform pricing somehow disproves that banks have market power. As far as we can tell, BS consider this the key insight of their broader focus on uniform pricing.

BS's theoretical argument is clearly incorrect. A bank can have significant market power and yet choose not to fully exploit local market characteristics at each branch. The definition of market power is that the bank can pay a deposit rate that is below the competitive market rate without losing its customers. This is exactly what banks do: paying deposit rates that are far below competitive market rates such as the Fed funds rate, the Treasury bill rate, or the yields of money market funds.

To be clear, this insight is well understood from work on non-financial firms. For example, DellaVigna and Gentzkow (QJE, 2019) find that it is common for retail firms to choose not to exploit differences in local pricing power due to concerns that this could harm their brand. Like non-financial firms, banks often do not necessarily set different rates at each branch. Banks are also subject to regulatory and legal constraints on deposit pricing. But make no mistake. Banks *do* exploit their market power at each branch by paying deposit rates that are well below competitive market interest rates. This is true at all branches, whether they are a ratesetter or not.

2.2.3 Market power and market concentration

BS appear to confuse market power and market concentration. They consistently seem to think that DSS argue that market concentration (HHI) is the sole determinant of market power.

DSS discuss at length that market concentration (HHI) is used for empirical identification to induce variation in market power. DSS also discuss that market concentration is only one determinant of market power. We reproduce the relevant discussion from DSS below:

Under the deposits channel, a bank's exposure to monetary policy depends on its market power. ***One important determinant of market power which we use for identification throughout the paper is market concentration. Yet as our model shows, banks also derive market power from other sources. These include product differentiation, the willingness of depositors to switch banks, their rate of participation in other markets, and depositors' financial sophistication and attentiveness.*** *A comprehensive measure of a bank's exposure to monetary policy under the deposits channel must account for all of these sources of market power. Our model provides a simple way to construct such a measure. A sufficient statistic for a bank's market power is the sensitivity of its deposit spread to the Fed funds rate, which we call the deposit spread beta."* (page 38).

Hence, DSS never claim that market concentration is the main determinant of market power. In fact, DSS propose a comprehensive measure of overall market power ("deposit beta") that is extensively analyzed and discussed in their paper. BS simply ignore this measure in their paper.

To be clear, it is a common feature of well-identified empirical work that the "exogenous" source of variation only explains some part of the observed variation. This characteristic applies to many papers, including DSS which uses market concentration to induce quasi-exogenous variation in market power. It does not follow that the quasi-exogenous variation has to fully explain the main variable of interest. This is true for DSS, as it is for the vast majority of well-identified reduced-form papers.

2.2.4 Deposit channel and large banks

BS claim that the DSS results do not hold for large banks. Again, this is incorrect. Rather, as it is the case for most papers analyzing banks, the empirical identification for large banks is challenging because the large bank sample is significantly smaller and the statistical power is limited. In fact, DSS show in the Internet Appendix Table IA.1 that they cannot reject the hypothesis that large banks and other banks have the same elasticity for rates and flows. BS ignore these results.

More importantly, DSS offer an alternative way to evaluate the deposits channel for large banks. Here is the relevant discussion from DSS:

"In order to have a significant aggregate effect, the deposits channel must affect large banks. Since the aggregate time series is dominated by large banks, it shows that large banks raise deposit spreads (Figure I) and contract deposit supply (Figure II) when the Fed funds rate rises, as predicted by the deposits channel. Yet, because monetary policy is endogenous, we cannot use the aggregate series to estimate its impact on large banks' lending. Instead, we again turn to the cross-section and reestimate the relationship between spread betas and flow betas for the subset of the 5% largest banks (this cutoff is commonly used in the banking literature; e.g., Kashyap and Stein 2000)." (page 40)

DSS show that their main results hold for large banks using the deposit beta. In making their assertions, BS simply ignore the evidence on large banks presented in the paper.

2.2.5 Evaluating the DSS research design

DSS provide a theoretical framework for the deposits channel and an empirical research design to estimate the deposits channel using aggregate, bank-level, and branch-level data. The empirical designs test the deposits channel using several different sources of variation:

1. Aggregate evidence: DSS show that deposit spreads increase and deposits flow out when monetary policy tightens. The effect is stronger for more liquid deposits. (Figures I and II).
2. Cross-sectional within-bank variation (using branch-level data): DSS show that branches operating in a less competitive environment increase deposit spreads more and have lower deposit growth when interest rates increase (Table II, Panel A, Col (1)–(3) and Table III, Col (1)–(3)).
3. Cross-sectional across-bank variation (using branch-level data): DSS show that banks operating in a less competitive environment increase deposit spreads more and have lower deposit growth when interest rates increase (Table II, Panel A, Col (4)–(6) and Table III, Col (4)–(6)).
4. Cross-sectional variation in expected monetary policy change (using branch-level data): DSS show that banks operating in a less competitive environment increase deposit spreads more when focusing on expected rate changes (DSS Table IV).

5. Cross-sectional variation in financial sophistication (using branch-level data): DSS show that banks dealing with less sophisticated depositors increase deposit spreads more and have lower deposit growth when interest rates increase (DSS Table V).

6. Cross-sectional bank variation (using bank-level call report data): DSS show that banks operating in a less competitive environment increase deposit spreads more and have lower deposit growth when monetary policy tightens (DSS Table VIII).

7. Cross-sectional bank variation using deposit betas (using bank-level data): DSS show that banks with higher deposit spread betas have lower deposit growth relative to banks with a lower deposit spread beta (DSS Table IX).

8. Cross-sectional bank variation in lending and real effects (using county-level data): DSS show that areas more exposed to the deposit channel experience less lending and have lower employment growth when interest rate increase (Tables VI and VII).

DSS argue that the deposits channel exists based on results from all these empirical tests. These tests use data at different levels of aggregation (aggregate, bank, branch), use different data sources (FRED, call reports, FDIC, Ratewatch), use different sources of variation, and address various identification concerns.

As is common in empirical identification, a subset of tests took the identification as far as possible by controlling for unobservable characteristics with a broad set of fixed effects. When doing this, a modern research design isolates variation that is well-identified, thereby often focusing on a relatively small part of the variation that is arguably exogenous. This is exactly the approach DSS take when using within-bank variation of deposit rates across branches of the same bank.

We believe that an assessment of the empirical research design needs to engage with the entire design and not just pick out item (2). BS fail to engage with most of the evidence presented in DSS beyond the within-branch regression columns in DSS Table II, Panel A, Col (1)–(3), and Table III, Col (1)–(3).

2.2.6 Funding substitution

BS claim that DSS aggregate results do not hold (see Table 11). They state that *"... there is no reliably negative relation between the change in log deposits and the Federal funds rate, either with or without lags."*

We are replicating Figure II from DSS. It seems obvious from the figures that there is a negative relationship between deposit growth and changes in the FF rate. It is hard to come up with a regression specification that does not find this result.

Figure 3: **Figure II from DSS**

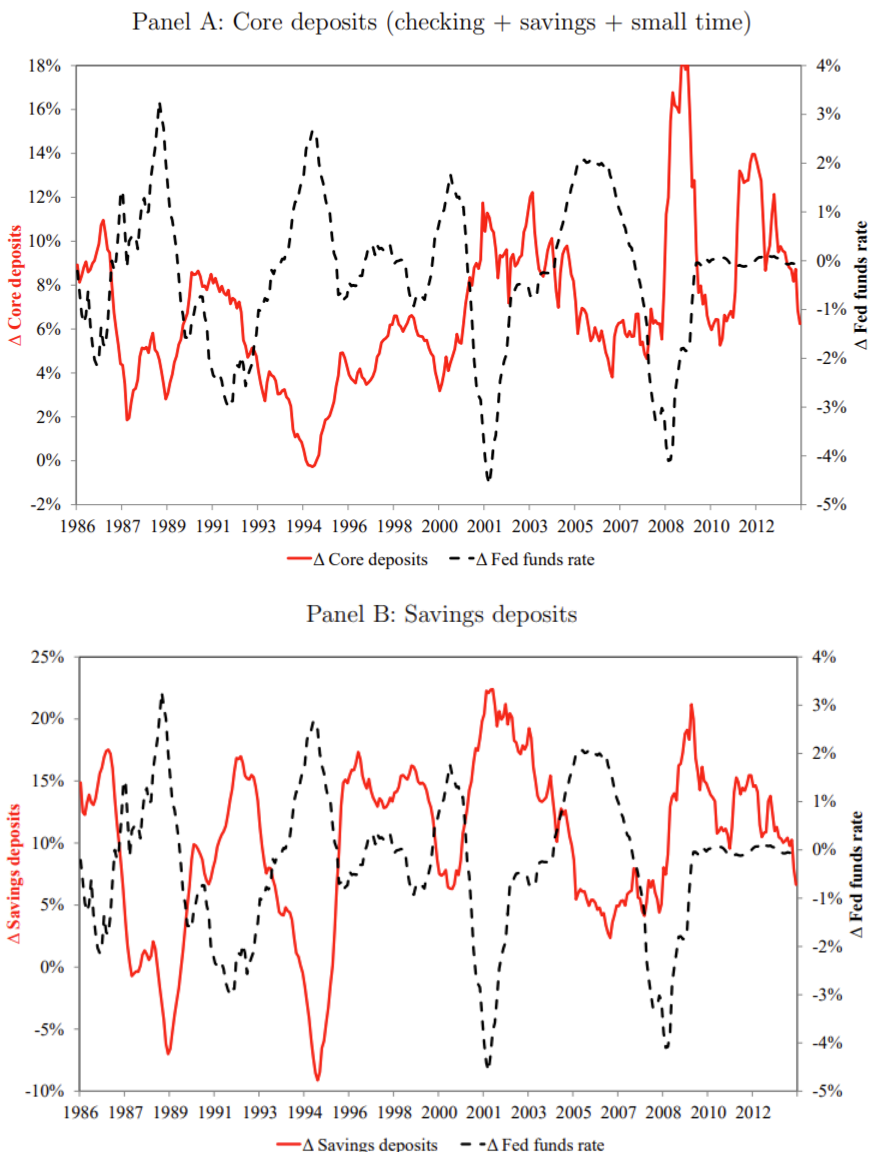

Moreover, BS arbitrarily exclude non-interest bearing deposits from their analysis. They claim that non-interest bearing deposits are not covered by the deposits channel. This is wrong. To the contrary, non-interest bearing deposits have a deposit beta of zero (by construction) and thus get significantly more expensive when the Fed funds rate increase. Thus they are very much subject to the deposits channel. Indeed, their zero beta is why banks invest significant resources in securing non-interest bearing deposits (see DSS 2021).

2.2.7 Bank-Level Deposit Channel Patterns

BS present some bank-level regressions in which they replace bank-level HHI with alternative measures of market power. DSS provide extensive results using the deposit beta as our preferred measure of bank market power. BS do not even replicate these results. Instead, they introduce some other measures that do not appear in our work. We suggest they start with the evidence in DSS.

2.2.8 Market Power measures and market valuation

BS claim that regression of market power on market equity divided by Tier 1 capital yields a insignificant coefficient. DSS never claim this regression would yield a significant coefficient. Market power does not come for free. In a competitive equilibrium, the cost of attracting and providing deposits equals its value. It follows that these regressions are not a proper test of market power. In any case, BS cannot rule out substantive market value effects given this result.

2.2.9 Uniform pricing in banking

BS never define what they mean by uniform pricing. If uniform pricing means that banks set the same deposit rate for all branches, it is straightforward to reject the claim. DSS uses within-bank variation and BS can replicate the original DSS results.

Alternatively, BS appear to use the term uniform pricing to mean that not all bank branches set different rates. This is a much weaker claim. The original BS abstract summarized this result stating that 85% of branches were networked branches.

This claim is not new. DSS already point this out in their paper: *"There is about one rate setting branch for every three non-rate setting branches in the data."* (page 1839). Moreover, there are at least two other papers that mention this fact. Both papers precedes BS by several years. Dlugosz et al. (forthcoming) study the real effects of local rate setting after natural disasters. Granja and Paixao (2019) study the impact of mergers on uniform rate setting. BS do not acknowledge that this point already appears in three earlier papers.

2.2.10 BS evidence on "heterogenous local demand"

BS explicitly acknowledge that they fail to provide an alternative theory that can account for the evidence in DSS. However, in some parts of the paper they claim that "heterogenous local demand" can perhaps explain the results. Specifically, they argue that *"[t]his leaves open the possibility that heterogeneous local demand shocks can be important drivers of deposit flows with implications for monetary policy transmission."* The main evidence in favor of this hypothesis is that county-level population growth (and county-level employment growth) is positively correlated with deposit growth. (see discussion on page 27 and Tables 13 and 14).

This claims falls short of an alternative explanation for several reasons. First, DSS never claim that local employment growth and population growth cannot be correlated with deposit growth. This correlation neither proves nor disproves DSS. Second, there is no attempt to explain the variation in the original DSS regressions. As discussed above, DSS provide a research design using several different identification strategies. The BS results on demand do not connect with any of them. Third, DSS show that price and quantity move in opposite directions using branch-level, bank-level, and aggregate data. This indicates that there was a shift to the supply curve, not a shift in the demand curve. If there are local shocks to demand, they would shift the demand curve (price and quantity moving in the same direction). Hence, it follows that demand shocks cannot explain the DSS results.

In any case, BS do not provide any empirical identification for their county-level results.

References

Begenau, J., & Stafford, E. (2023). Uniform rate setting and the deposit channel. *Available at SSRN 4136858.*

DellaVigna, S., & Gentzkow, M. (2019). Uniform pricing in US retail chains. *The Quarterly Journal of Economics, 134(4)*, 2011-2084.

Dlugosz, J., Gam, Y. K., Gopalan, R., & Skrastins, J. (Forthcoming). Decision-making delegation in banks. *Management Science.*

Drechsler, I., Savov, A., & Schnabl, P. (2017). The deposits channel of monetary policy. *The Quarterly Journal of Economics, 132(4)*, 1819-1876.

Drechsler, I., Savov, A., & Schnabl, P. (2021). Banking on deposits: Maturity transformation without interest rate risk. *The Journal of Finance, 76(3)*, 1091-1143.

Granja, J., & Paixao, N. (2023). Bank consolidation and uniform pricing. *Available at SSRN 3488035.*