What's In Their Wallet?

Jonathan Sandberg
Princetonbankconsortium.com

Abstract:
Why is the bank Capital One's 2015 NIM such an outlier relative to Bank Holding Companies (BHCs) of comparable and larger asset base? Capital One NIM is double the average bank NIM in 2016, Why is that? Press coverage over that past couple years attribute COF outperformance to going long risky credits and managing the write-downs. Maybe they are right and COF is just good at making and managing loans for a decade. On the other hand, maybe COF is running a very different quantitative model than its competitors. COF has 300+ BN of assets in 2016 and is making 300 more bps per year on their assets than the average competitor BHC makes on their assets. Banks' NIM are at 30-year lows but not at COF. Maybe the competitors can learn from COF. We discuss some of the quantitative possibilities.

Net Interest Margin

This note is divided into three sections. This first section, *Net Interest Margin*, discusses the Net Interest Margin current status both in the US market and from a quantitative/technology modeling perspective. The second section, *Capital One Financial Net Interest Margin*, quantifies the comparative advantage of the elevated COF NIM relative to the recent COF earnings disclosures and discusses the analyst and media coverage of COF. The final section, *COF NIM The Technology Case*, outlines the case for assuming Capital One is running some form of stochastic accrual portfolio simulation perhaps topped by a LP/NLP optimizer guided implementation of the capital allocation plan.

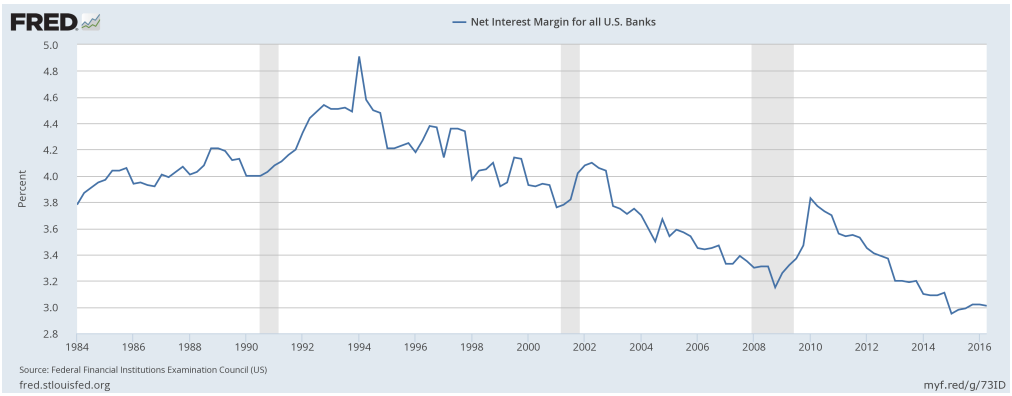

Figure 1: Average Bank Holding Company (BHC) NIM

Bank NIM is down and interest rates are expected to stay down (see Rieder). Figure 1 shows the St. Louis Federal Reserve figures for the average Bank Holding

Company (BHC) NIM over the past 30 years. Notice the current average NIM is at a 30-year low.
Capital One's NIM, on the other hand, is not at a 30 year low. In fact they buy assets on the market from ING, GE, and Chevy Chase yet still maintain their historic high NIM levels. COF assets grew 8x from 2004 to 2015 yet the COF NIM level stayed in a narrow range, occasionally busting out topside. How do you do that? We will review the evidence that the correct figurative question is not "What's in their Wallet" but "How do they choose what's in their Wallet?" We will make the case that COF is running some sort of novel Dynamic Stochastic Optimization for implementing their capital allocation plan. COF's Wallet is filled with securities selected through an LP/NLP optimization process and possibly some Dynamic Programming feedback control process. That is how they make 300 bps more than mostly every other bank in a punishingly low interest rate environment. Technically this appears to be more of an internal efficiency game than a zero sum game. In 2016, COF has a free hand in this game, because no one else knows how to play enough to move the dial.

Here are the steps to play the game of contemporary ALM Management NIM Optimization. There are several vendors who can help banks with parts of the game infrastructure including QRM, Bancware, Polypaths, Kamakura, Oracle, and IBM. Some banks, like COF, are just further along in implementing the necessary business process. The ideas are straightforward but the implementation can be non-trivial.

1. Marking the Accrual Portfolio during capital plan implementation,
2. Static Simulation of the Accrual Portfolio,
3. Stochastic Simulation of the Accrual Portfolio,
4. Stochastic Optimization of the Accrual Portfolio from Static Start Date, and
5. Dynamic Stochastic Optimization of the Accrual Portfolio

These steps breakdown to the following questions:

1. Do you have all the securities in the Accrual Portfolio for today? A large bank could have upwards of a couple hundred million securities in their accrual portfolio. This is the first order factor determining the size of the quantization error in NIM optimization. The quantization error impacts the optimization a couple ways. First it allows the theoretically optimal NIM to diverge from the realized actual NIM in matching assets and liabilities. Second, it obscures the delays in timing and shortfall/excess quantity in the simulation of the capital plan implementation. Quantization error arises in the enforcement of the optimization constraint that the asset notional is equal to the liability notional at the end of each business day. If it takes 12 weeks to find the assets and liabilities then Treasury operation may find during the quarter that they need to borrow to cover excess assets or that they have missed opportunities with excess liabilities. In either case, the realized NIM differs from the computed/optimized NIM. If you only know your Accrual Portfolio cashflow positions every quarter, you will have to live

with relatively large cashflow arrival time errors (6 weeks off perhaps on average).

2. Do you have accurate security level cash flow models for each security in the Accrual Portfolio and each new business security? You need a cashflow model that handles, interest rate, credit, and FX risk. The model attribution of default loss and loan impairment is useful in making the optimization converge to a realistic maximum NIM and facilitates P&L attribution of realized versus theoretical NIM back down to the security level. The computational form of the model is important for the overall system performance. On contemporary x86 vector processors (read AVX2) you need about 7 or 8 picoseconds per floating point multiplication or addition if the computational form of the model is simple expression evaluation.

3. Do you have a stochastic market model covering all the econometric and market data variables needed by the Accrual Portfolio security models? You need a stochastic market model for the Monte Carlo simulation random number generation. The Monte Carlo simulation allows the optimization process to know the expected distribution of cash flow receipts and payments through an accrual portfolio forecasting time period.

4. Do you have enough FP (floating point arithmetic) processing power to stochastically find the expected cashflow distribution over time and then run an LP/NLP optimization to find the optimal capital plan implementation? Note you need to take advantage of the fact that credit returns on individual securities can cause losses in subsequent simulation periods. You do that by choosing security cash flow models that handle interest rate, credit and fx risk accurately. Also you need to observe that the bulk of the accrual portfolio is buy and hold and hence constant with respect to the LP/NLP optimization. This is a good thing because the expected runtime complexity of the LP/NLP optimization is cubic in the number of security choices. The system runtime will be far better optimizing the choice of new investments over a given period (e.g., couple of hours and a contemporary server) than trading in an out of potentially any of the 100m securities in the accrual portfolio.

5. Finally do you have enough processing power to take dynamic feedback from the market or the capital plan implementation to reoptimize based on the feedback information?

My rank speculation, using public information, is that COF has a production implementation of #4. Stochastic Optimization with aggregated security positions and good security models but maybe cannot quite get to daily simulation periods. Possibly the rate limiting factor is some combination of the security models and the Monte Carlo simulation execution runtime, even with large amounts of compute capacity. QRM sounds like a C# shop and that could do it for performance if they are

supplying the stochastic Monte Carlo simulation code. Let's look at the data to see what fits.

Figure 2 shows the NIM published in the quarterly and annual reports of Capital One (COF), JP Morgan (JPM), Wells Fargo (WFC) and US Bancorp. COF NIM enjoys a 200 to 300 basis point advantage over these competitors.

JPM NIM is probably 15 to 20 bps depressed to cover the London Whale losses in 2012 -3. These positions are interesting because JPM CIO was ostensibly using the Whale's positions to hedge the Corporate Credit risk in the accrual book. That in turn means the Treasuries at the various banks will put on hedges (in size) for the accrual book to manage the aggregated risk. So the hedging cost will be an important part of the capital plan implementation.

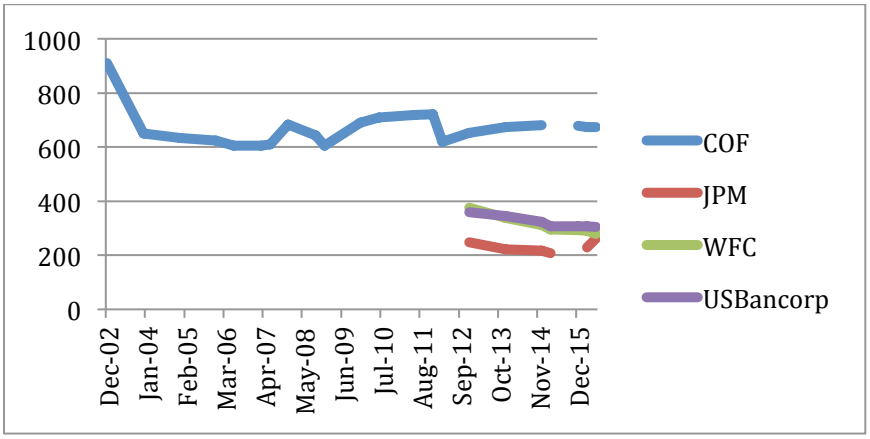

Figure 2: NIM for COF, JPM, WFC, and US Bancorp.

JPM NIM appears to be converging with WFC and US Bancorp in 2016. But the overall trend of NIM compression 2012 to 2016 appears to be shared by these three banks. COF NIM is not compressing in any significant way in the period between 2012 and 2016. In fact it is reasonably flat. The COF dip in NIM in 2008 is understandable in the context of the loan write-downs and displacement post the credit crisis as well as the Chevy Chase purchase. The dip in 2011-2 is puzzling. The ING Direct purchase was in 2012 and COF put in $6.2b cash to the deal. That could do it. Apart from acquisitions, the NIM behavior for COF is remarkably different than its competitors and the average BHC NIM shown in Figure 1 from 2010 onward.

Figure 3 shows the reported assets at each of the banks COF, JPM, WFC, and USBancorp. Maybe COF has a small asset base size advantage over WFC and JPM. In 2016 WFC and JPM have approximately $2tn of assets, US Bancorp ~$438bn assets, and COF ~$339bn assets.

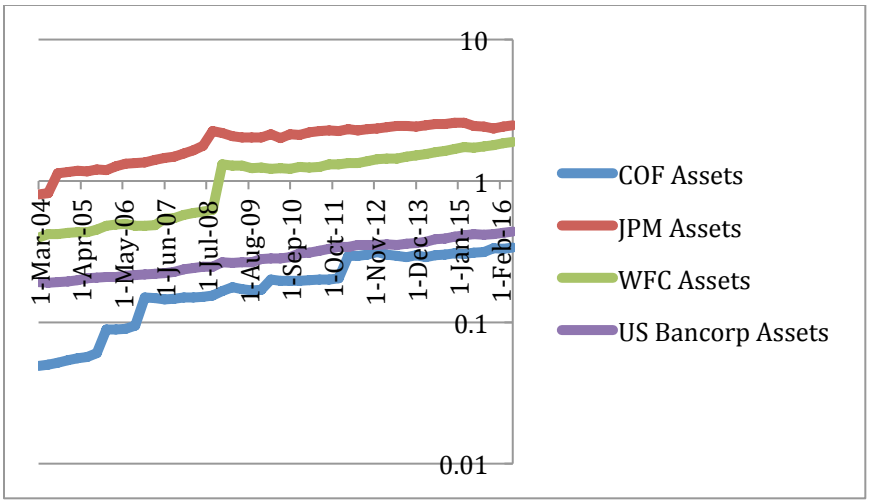

Figure 3: Bank Assets in Trillions USD from 2004 to 2015.

Look at the recent asset growth rates of JPM, WFC and US Bancorp. You can see where Wells purchased Wachovia in 2008. Even though US Bancorp is much smaller than JPM and WFC the asset growth rates are similar. If you refer back to Figure 2 you can see the NIM of WFC and US Bancorp are very similar (and dominate JPM over the past several years) despite the nearly 7x asset size difference. So US Bancorp doesn't show any particularly notable advantage for having only $438bn in asset versus nearly $2tn in assets for WFC. Maybe COF is doing something to jump the NIM that could work in scale. Notice that COF has been ready to buy assets: ING, GE Healthcare, and Chevy Chase being notable examples.

Figure 4 shows selected historical market data from the Baseline Supervisory Dodd-Frank Stress Scenario. We have chosen all market data prior to Jan 2016 assuming that they are accurate historical levels.

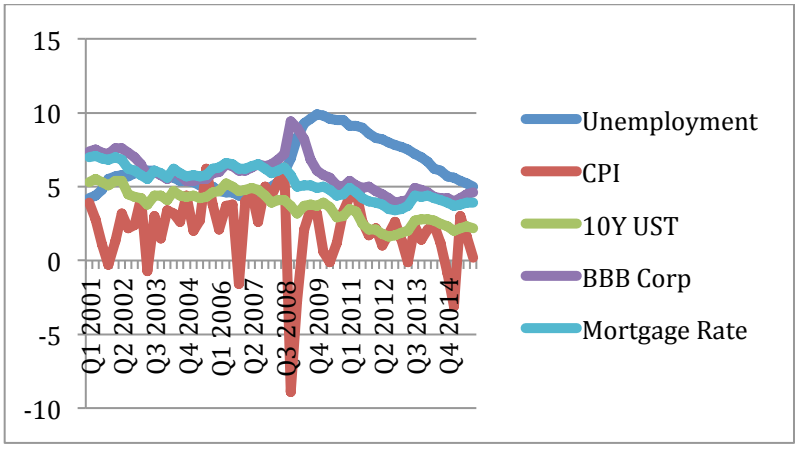

Figure 4: Baseline Dodd-Frank Historical Market Levels.

Short term UST yields are very low. Funding through either deposits or wholesale is extraordinarily cheap. BBB Corporate spreads are down from their peaks in 2009. Unemployment is back at pre-crisis levels. If COF had good security cash-flow

models the UST, BBB, Mortgage Rate, and CPI input variables where all reasonably well behaved from 2010 onward. The unemployment numbers could have been a surprise in the period after 2010. Other than that, the underlying market and econometric variables seem relatively stable.

Finally figure 5 shows the US Fed data on loan impairment and default historical write-off levels. COF NIM should have been hit in the period 2010 to 2013 like all its competitors (see Figures 1 and 2) but apart from the ING Direct purchase nothing much changed in the COF NIM level of growth.

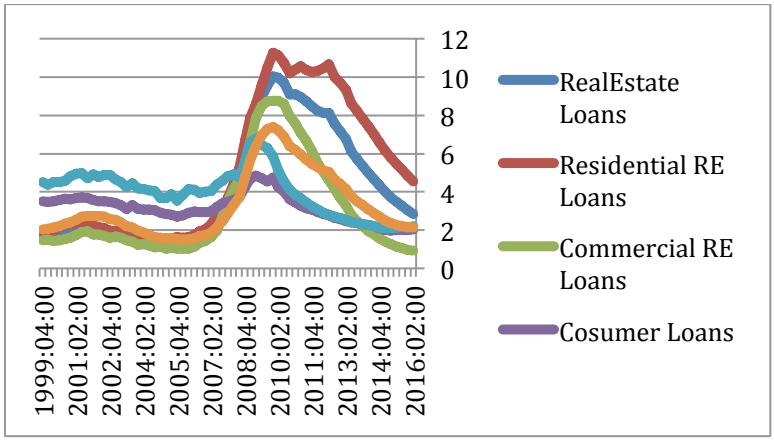

Figure 5: Loan Impairment and Default Writedown levels 2000 to 2016 in percent.

The Liability side of ALM has not been a driver of NIM with rates so low. NIM has been driven by asset selection, hedging, and regulatory rollout. It sure looks like COF NIM is getting managed very consistently and differently from the competition.

Capital One Financial Net Interest Margin

COF generates an additional $10bn of annual revenue with the 300+bps NIM margin it maintains over the FRED average BHC NIM (see Figure 1). The July 2016 review of COF Second Quarter results shows $12.5bn Total net revenue and Provision for credit losses $3.1bn for the prior 6 months. The NIM optimization margin is generating approximately $5bn of that revenue. It may be debatable how much of the incremental loss reserve should be attributed to the NIM optimization but let's say it is $1.6bn The technology running the NIM optimization, at 36% of the semiannual Total net revenue (3.4b/9.4b), is a dominant feature of the COF financial results. COF may no longer be steering a rocket with a joystick. It is like Uber for assets.

The NIM optimization technology financial results flies just under the radar in the equity analyst and press coverage of COF. Morgan Stanly Equity Analyst covering COF, Betsy Graseck notes in Nov. 2015 Brookings Institute Q&A:

DR. BERNANKE: Humility is always good to have.
Betsy, you were talking about how the LCR would push banks into riskier investments. In principle, assuming the constraint on the capital requirements, the risk weighted assets, and stress tests, assuming that's the case, in principle, and obviously the risk weights might not be right and so on, but in principle, wouldn't that solve the problem, because you need more capital for the riskier investments?

MS. GRASECK: Yes, that's accurate, but you also are trying to triangulate to a certain outcome, and as a result, if your RWA is your binding constraint, you have a little bit less wiggle room than if it's your SLRs, your binding constraint.

To your point, it depends on which one is your binding constraint. Institutions migrate between them. At the same time -- I guess that's part of the challenge with the all the different rules, at the same time they are trying to triangulate as a goal to at least earn their cost of equity.

Saul Perez analyzed COF NIM in 2015 in Market Realist, Why Capital One's net interest margin is above the sector average. Perez observes

"In an environment of falling net interest margins, Capital One has been able to hold on to its margins. At the end of 2014, the net interest margin was at 6.17%. This was almost double the sector-wide average net interest margins of 2.91%. This is due to Capital One's presence in high net interest margin loans.
Wells Fargo (WFC), U.S. Bank (USB), and PNC Bank (PNC) have much lower net interest margins than Capital One. Capital One has one of the highest net interest margins among the banks in the Financial Select Sector SPDR (XLF)."

COF appears to have determined how to "triangulate" as Betsy Graseck puts it and Optimize the Net Interest Margin, subject to various capital constraints. Perez clearly highlights the historical disparity of COF NIM versus competitors. One remaining question is how does COF "triangulate" and maintain industry leading NIM year over year?

Let me offer an alternative explanation for why COF NIM stands out so singularly among BHCs between 2010 and 2016. COF is running a more advanced NIM optimization of their capital plan implementation than any of their competitors. The evidence presented here is circumstantial. I do not have any direct confirmation from anyone currently at COF or anyone previously employed or partnering with COF. The technology case is very plausible.

COF NIM the Technology Case

Let's outline the case for COF running some form of NIM optimization, If you read around for the positions COF is trying to fill they ask for candidates with experience in QRM. I don't have any particular experience with QRM but I can read the spec sheets and get the gist of what's going on. Seems likely that QRM, Bancware, Polypaths, Kamakura, Oracle, and IBM offer roughly the same services to help banks develop code to run stochastic simulations and possibly LP or NLP optimization on top of their capital allocation plan. I do not know for sure, but I would bet, there are

other banks that use QRM for ALM, so QRM alone is not the secret sauce delivering COF's decade long outstanding NIM performance.

COF probably has between 1mm and 10mm securities/accounts in their accrual portfolio - let's call it 10mm contracts total to be conservative. Let's assume they aggregate securities/contracts on average 100 to 1 so there are 100K positions to track daily through their reasonably sophisticated security models. Let's assume they can get daily, or at least weekly, full accrual inventory, with new business, globally in production. That puts them ahead of several larger banks. They know MBS and ABS security cashflow/prepayment models and have production code for them. Perhaps they use YieldBook and Intex or have their own proprietary code. They probably have a stochastic market model for their security models. The various vendors have market models for sale. Getting standard market and econometric data is easy post CCAR. I would expect they are targeting getting to 10K Monte Carlo paths and daily simulation intervals with the vendor Monte Carlo code. The coding performance issue typically involves integrating the security models with the Monte Carlo simulator so the FP execution pipelines stay full and the memory hierarchy does not suffer many non-compulsory misses. For static simulation and optimization the vendor code works just fine, given enough cores.

Figure 6 lists some of the projected Monte Carlo simulation performance levels single core. We show the projected single 4GHz core execution time depending on the average processor cycles required per security, the number of securities in the accrual portfolio, the number of time intervals in the Monte Carlo simulation, the number of paths required by the Monte Carlo to get a converged expectation for cash flow receipt and payment, and the aggregate Monte Carlo simulation time scaled to seconds, minutes, hours, and days.

Stochastic Simulation	COF Theoretical	COF exp	COF opt	JPM exp	all BHC ALM
average cyles per security model.	100	10,000	20	10,000	20
Number of securities	100,000	100,000	100,000	1,000,000	1,000,000,000
number of time intervals	1250	60	1250	60	1250
number of paths	10,000	10,000	10,000	300	10,000
Stochastic Simulation time at 4GHz	31,250	150,000	6,250	45,000	62,500,000
sec	31,250	150,000	6,250	45,000	62,500,000
min	521	2,500	104	750	1,041,667
hrs	9	42	2	13	17,361
days	0	2	0	1	723

Figure 6: Expected NIM Optimization Runtimes - single core

The performance is listed for several scenarios: COF Theoretical, COF expected, COF optimal, JPM expected, and All Bank Hold Company ALM optimal. Without confirming data we see COF's stochastic simulation performance falling somewhere between COF Theoretical and COF expected. You can do something like synthetic variables to perhaps half the required Monte Carlo paths to get convergence. The main performance variables that you can control, however, are the number of time intervals and the average processor cycles per security model. The problem with dialing down the number of time intervals is the increase in the simulation

discretization error. Other than massive parallelization, the main parameter to control the simulation performance is the average cycles per security model. The good news is that is typically quite feasible. The bad news is that not many programmers know how to do that on commodity microprocessors. The opportunity as you can see in the last scenario is that if you control the average cycles per security level then you can reasonably run all the US BHC assets and liabilities with daily time intervals out 5 years on a single core in 723 days. First observation is that can be run in a parallel a dozen or so Intel Knights platforms for a very modest hardware cost. Second observation, if you are acquiring assets it is probably good thing to know what the stochastically simulated value of those assets are to you as well as the competition. The last scenario addresses the computational cost of that competitive feature. Certainly, you (if you are not the US FED) don't necessarily know all the individual securities in the accrual portfolios of all the BHCs, conversely you do know the aggregate security distribution over time.

Current competitive Linear Programming solvers are good now for upwards of 1mm variables in about an hour of execution time on standard benchmarks running on a vanilla x86 server (see Mittlemann). Of course this also depends on the vendor software selected. Once you know the expected distribution of the accrual portfolio cash flows you can use the LP and the pool of new business securities to optimize the NIM subject to regulatory and business constraints. You can represent the capital allocation plan through Transfer Pricing curves and target branch efficiency or just outright. One good part for execution time is the combinatorial search for the maximum NIM does not depend on the number of securities in the accrual portfolio. That stochastic simulation runtime cost is a one-time execution.

I think COF generates their extra 300+ bps of NIM from the LP/NLP solvers finding:

1. Timing and size of new business purchases
2. Efficiency of Capital plan implementation at 600 or so branches and subs
3. Statically rerunning the capital plan optimization post significant market changes/news

That would fit a production runtime of between 9 and 42 hours on a single core. Assuming it can be parallelized to some degree, without any heroics, puts COF's production runtime somewhere between 10 minutes and the better part of a day (see Figure 6).

The funding situation currently is rather simplified because rates are so low. The places to make money are in new business selection and efficient hedging as well as developing a model of the branches/subs efficiency in a capital allocation plan implementation. Without access to the actual data, I would expect the model for the capital allocation plan implementation is the dominant money-maker currently. But that will change over time. That is why this NIM Optimization process in not necessarily a zero sum game.

Summary:

We review the circumstantial evidence that in the period 2010 to 2016 COF numerically optimized their capital purchases to maximize Net Interest Margin.

If we are seeing this accurately, COF is going to buy assets and deposits to take advantage of its NIM stochastic optimization capabilities. It would be reasonable for them to expand their asset coverage to include M&A targets so COF knows what the expected cash flows are worth to them as well as estimating the value of the cash flows to the seller. I suspect that they can do this and maybe used it for the 2012 ING purchase?

References:

Bancware, SunGuard. here.

Board of Governors of the Federal Reserve, Charge-Off and Delinquency Rates on Loans and Lease at Commercial Banks, here.

Board of Governors of the Federal Reserve, Jan. 2016, 2016 Supervisory Scenarios for Annual Stress Tests Required under the Dodd-Frank Act Stress Testing Rules and the Capital Plan Rule, here.

Brookings Institution, Nov. 2015, Are we Safer? A Look at the Financial System, Post-Crisis, here.

Capital One, Press Release, Capital One to Acquire ING Direct for $9.0 Billion in Stock and Cash, here.

Federal Reserve Bank of St. Louis, Net Interest Margin for all U.S. Banks, here.

IBM, Algo Asset Liability Management, here.

Mittlemann, Decision Tree for Optimization Software, here.

Morgan Stanley & Oliver Wyman, Wholesale Banks & Asset Managers, here.

Oracle Financial Services, Asset Liability Management: An Overview, here.

Saul Perez, Market Realist, Why Capital One's net interest margin is above the sector average here.

QRM, Quantitative Risk Management, here. QRM Analytical framework described as ETL from SQL Server, tools for managing parallel jobs, and OLAP for viewing the results. Not much prose spent on the computational simulation or Optimization calculation. The Behavioral Model Research is interesting:

"This group works with clients to develop specialized behavioral models in order to estimate how the customers of a client will respond to changes in market and economic conditions. Among the types of modeling QRM assists with are models for prepayments, fallout, deposits, and credit events."

QRM calls it option-adjusted deposit valuation model in their write up, here. I assume this means a cash flow model that handles the depositors' option to draw down their deposited cash.

Rieder, BlackRock, Another reason rates are likely to be lower for longer, here.

Julie Segal, Institutional Investor, 5 Years after Dodd-Frank, U.S. Banks Dominate More Than Ever, here.