BPS HAWK

bpshawk.com

Why Your Clients Need This

The Broker's Guide to Mortgage Rate Monitoring

Backed by peer-reviewed research from leading economics journals

What's Inside

BPS Hawk

Why Your Clients Need This

The Broker's Guide to Mortgage Rate Monitoring

> **The short version:** To know when a client should refinance or buy points, you need two things: a rule tailored to their situation, and something that checks rates against that rule on their behalf. Those are the two problems BPS Hawk solves. The research shows that between the two, attention is far more important than which rule you use. The rest of this document explains how and why.

> **Bottom line:** Every mortgage your client takes out comes with a built-in refinancing option worth 4-4.6% of the loan balance. On a $500,000 mortgage, that's over $22,000. Research across 1.8 million mortgages shows that the average borrower loses at least 20% of that value, roughly $4,000-$5,000, purely from not paying attention. With active monitoring, nearly all of it can be captured.

Backed by peer-reviewed research from the Journal of Money, Credit and Banking, American Economic Review: Insights, Journal of Financial Economics, Management Science, and Real Estate Economics.

If you already have a break-even rule you're comfortable with (e.g., 36 months), you can use that as your target in BPS Hawk. The system will monitor your clients against whatever target you set. You can also see how your rule compares to the research-backed suggestion on each client's sensitivity table. The most important thing is not which rule you use. It's that you're attentive. Don't get bogged down in the math. The system handles it.

Author's Note, Daniel Fernandez

The #1 thing I'm proposing is this: set a rule ahead of time and follow it. If you already have a break-even rule (24 months, 36 months, whatever you're comfortable with), your rule is probably not that far off from what the research says is optimal. The ADL model shows that different reasonable rules produce similar option values at origination. So the rule matters less than you'd think. What matters far more is that you're attentive; that you're actually watching and acting when the math says it's time. Everything else in this document (the models, the formulas, the research) is there to back that up and quantify it. But the core idea is simple, and I'll explain it below.

I don't have a PhD in economics. I want to be upfront about that, because the research in this document draws on peer-reviewed academic work that I'm interpreting and applying, not producing myself. My analysis is not peer reviewed. Take everything here with a grain of salt, and I encourage you to look into the papers yourself. The citations are at the end of this document. A lot of what follows (the industry commentary, the framing, the conclusions I draw) is my opinion. But what I'm actually proposing is straightforward: set a rule and follow it.

How I got here: I started thinking about this in November of 2025. I was reading online blogs about mortgage refinancing and kept running into rules of thumb that didn't seem rigorous. That led me down into the academic research. In reaching out to academics, it became clear that the ADL model (Agarwal, Driscoll & Laibson, 2013) was the one with a consensus behind it. It's the standard reference for optimal refinancing thresholds, published in the *Journal of Money, Credit and Banking* and cited across the field. The Berger extension (2025) builds on the same model and adjusts the threshold downward for borrower inattention. Both are peer-reviewed and well-established.

At first, I was very focused on finding the right rule. Then I found Berger's research on inattention and realized that people are far less attentive than I had assumed, evaluating refinancing once every 3-4 years on average. That makes logical sense when you think about it, and it shifted my focus entirely. The rule matters, but attentiveness is the much bigger lever. This is the main finding I'm riding on, and it's most clearly demonstrated in Berger et al.'s (2024) paper *"Refinancing Frictions, Mortgage Pricing and Redistribution"*, which shows that inattention is a massive friction and that the slowest borrowers overpay by 60 basis points per year compared to the fastest.

Why this matters to me personally: My mom is a loan officer with over 20 years of experience. It became clear to me early on that borrowers don't have the right framework and LOs don't have the tools to give them one.

I also believe that the large, impersonal mortgage corporations operate in a way that is fundamentally at odds with what I'm proposing here. Everything should be thought of as client-first and worked backward from there. These companies have not built anything like this because they do not want transparency. Their model is to sell the client a loan, collect the fee, and move on.

The academics know there is a consensus around this research. The ADL model is well-established, the inattention findings are robust across multiple studies, and the implications are clear. Yet none of the major players talk about any of it. That is either willful deception or blissful ignorance, because their methodology of selling a client at any cost does not allow them to take a client-first framework when approaching this question.

Borrowers are constantly anxious about rates ("should I wait? should I lock? rates might drop more?") and they're bogged down by the news constantly. They cannot have peace of mind, and the mortgage broker loses control to steer them in the right direction. If borrowers were explained all of this ahead of time (the option value, the threshold, the plan), they could save real money and stop worrying.

This is especially clear at purchase. It's very common for a borrower to say "I'm getting X rate now, but I'll refinance later." They already know this. By addressing it upfront with a specific, research-backed number ("your refinancing option is worth $22,000, and here's the threshold where it makes sense to act"), the broker shows a much higher level of sophistication and builds trust from day one. The decision is already made before the opportunity arrives.

What I built and why: I realized early on that I needed to boil everything down to one number. The ADL model produces a threshold in basis points. The full break-even with amortization, taxes, MI, and all the adjustments can tell you the day it becomes positive, but it still doesn't tell you which

rule is optimal. The welfare value from ADL shows that the rule doesn't matter that much (different rules produce similar option values at origination). So rather than force LOs to understand basis points and option theory, I converted the suggestion to a break-even in months. One number the client can commit to, consistent across the rate table, easy to explain.

I also needed to solve the pricing problem without relying on APIs, wholesalers, or services like LoanSifter. So I built a system where base pricing grids are uploaded daily. In my research, rates at each level don't vary significantly on a given day, so this approach is accurate without the cost or complexity of live API feeds. LoanSifter charges around $100/month for their API, and Banking Bridge starts at $300/month for live rates. Neither of them offers threshold monitoring or break-even suggestions. For now, BPS Hawk works for Conventional and FHA loans.

What I'm proposing: The #1 thing I'm advocating is simple: set a rule and follow it. In order to follow it, you need a system that monitors for you, because doing it manually across your entire book is either impossible or an enormous time sink. I don't believe most LOs are doing this effectively, even the ones who say they track rates for their clients. And there's nothing else out there that does it.

I believe this is the biggest lever in mortgage today. It's not just about making more money as an LO. BPS Hawk is 100% focused on the consumer. What other system could you implement that saves the average borrower roughly $4,000-$5,000 in present value terms at origination? LOs have a competitive incentive to do this, but I also believe they have a moral obligation. Borrowers are leaving thousands of dollars on the table because nobody is watching.

I look forward to building this out, meeting loan officers, and using the data we collect over time to do further research on how borrowers actually behave on our platform. As we gather real-world data, the estimates in this document will get sharper. For now, set a rule. Follow it. Be attentive. That's the biggest thing you can do for your clients.

Daniel Fernandez, Founder, BPS Hawk

---

1. Every Mortgage Has a Built-In Option Worth Real Money

When a client takes out a mortgage, they don't just get a loan. They get the right to refinance if rates drop. That right has a real, quantifiable dollar value, just like an insurance policy they received for free at closing.

Agarwal, Driscoll & Laibson (2013) calculated exactly what that option is worth:

Mortgage Size	Option Value	% of Balance
$250,000	$9,859	3.94%
$500,000	$22,244	4.45%
$1,000,000	$47,531	4.75%

Source: ADL (2013), Tables 5-6, p. 609. Option value equals the welfare loss from using a simple break-even rule, which ignores the option entirely.

What does this mean in plain English?

A $500,000 mortgage has over $22,000 in value sitting inside it at origination. That value can only be captured by refinancing at the right time. If they're not watching, that value expires like an uncashed check.

The option value is approximately 4-4.6% of the loan balance for mortgages in the $250K-$1M range. This value exists from the moment the loan is originated. It's built into the mortgage.

---

2. How Much Are Borrowers Losing?

The research is consistent across multiple studies: the average borrower loses a significant fraction of their refinancing option value, and the primary cause is inattention, not the wrong rule. No single study gives us a clean "this is exactly how much inattention costs per borrower." The data comes from different angles, time periods, and methodologies. But taken together, the picture is clear. Here's what each study found:

Study 1: Keys, Pope & Pope (2016), "Failure to Refinance"

Journal of Financial Economics. 376,036 mortgages. Data: CoreLogic, covering ~85% of the U.S. mortgage market.

What they looked at: As of December 2010, when 30-year mortgage rates were at 4.3%, they identified borrowers who were qualified to refinance (FICO >680, LTV <90, never missed a payment, current CLTV <90) and for whom refinancing was clearly in their financial interest, and asked: how many hadn't done it?

What they found:

- **20% of these qualified, good-standing borrowers had not refinanced** despite it being clearly optimal.
- The median present-discounted loss for these non-refinancers: **$11,568 per household**.
- The unadjusted lifetime loss: **$45,473 median**, with 11% losing more than **$100,000**.
- Nationwide, they estimated the total forgone savings at **$5.4 billion per year**.
- Two years later, **40% of those same non-refinancers still hadn't acted**, even as rates dropped further.

Important context: This 20% is a snapshot at a specific point in time, during a particularly favorable rate environment (rates had been falling for years after the 2008 crisis). In a different rate environment, the percentage of borrowers for whom refinancing is "clearly optimal" would be different. But the behavioral finding, that a large fraction of qualified borrowers simply don't act even when the math is obvious, is what matters.

The behavioral evidence: A nonprofit mailed 446 pre-screened, eligible borrowers a pre-approved refinance offer with zero upfront costs. Only 16% responded. Of the non-responders surveyed: 25% didn't even open the letter, and 33% said they planned to call the loan officer but never got around to it. These borrowers didn't fail to qualify. They failed to act.

Study 2: Agarwal, Rosen & Yao (2016), "Why Do Borrowers Make Mortgage Refinancing Mistakes?"

Management Science. 271,216 refinancing observations. Data: Government-sponsored enterprises, covering ~25% of national mortgage transactions, 1998-2011.

What they looked at: Among borrowers who DID refinance, how well did they time it? Did they refinance at the right rate, or did they act too early, too late, or at the wrong rate?

What they found:

- **57% of borrowers refinance suboptimally**, making at least one type of error.
- **50% choose the wrong rate**, refinancing on average **60 basis points higher** than the optimal rate for their loan.
- **17% wait at least 6 months too long** after the opportunity appears.
- The average borrower has **2.66 months of "burnout"**; the opportunity sitting there before they finally act.
- Brokered and non-brokered refinancings show **identical error patterns**. Without a monitoring system, brokers are guessing just like everyone else.
- Second-time refinancers make smaller errors, **evidence of learning**. Borrowers who've been through it once do better the second time.

Study 3: Berger, Milbradt, Tourre & Vavra (2025), "Optimal Mortgage Refinancing with Inattention"

AER: Insights. 780,000 Fannie Mae mortgages.

What they looked at: How often do borrowers actually evaluate whether to refinance, and how does that inattention affect their behavior?

What they found: The average borrower evaluates refinancing about **once every 4.3 years**. This isn't directly surveyed; it's backed out from the data. They measured the average rate gap across all borrowers (0.15%) vs. the average rate gap at the moment of prepayment (0.78%). The difference (0.63%) implies an attention rate of 0.234/year, meaning the typical borrower seriously evaluates refinancing about once every 4.3 years.

A note on what "evaluating" means: this isn't glancing at a headline or hearing about rates from a friend. It means actually sitting down and assessing whether refinancing makes sense for their specific loan, looking at their rate vs. current rates, considering costs, and making a decision. Most people see rates in passing but never do the math.

Study 4: Berger, Milbradt, Tourre & Vavra (2024), "Refinancing Frictions, Mortgage Pricing and Redistribution"

NBER Working Paper. 250,000 Equifax CRISM borrowers, 2005-2017.

What they looked at: The same authors in a separate study estimated borrower attention rates from a different dataset and examined how differences in refinancing speed create cross-subsidies between borrowers.

What they found:

- **45% of borrowers almost never evaluate refinancing** (effectively zero attention rate).
- Population average attention rate: **once every 3.3 years**, consistent with the 2025 study.
- **84-92% of borrowers are "asleep" in any given quarter** (citing Andersen et al., 2020).
- The slowest refinancers overpay by **60 basis points per year** compared to the average. The fastest refinancers pay **121 basis points per year less**.
- On a $500,000 mortgage, that's a difference of **$3,000/year** for the slow borrowers.
- These cross-subsidies are **regressive**. Lower-income, lower-FICO, and younger borrowers tend to be less attentive and end up subsidizing wealthier, higher-FICO borrowers who refinance quickly.

Our estimate: at least 20% of the option value is lost to inattention

None of these studies gives us a single, clean dollar figure for what inattention costs the average borrower. The data comes from different time periods, different rate environments, and different methodologies. Calculating the exact loss would require knowing future rate movements, which no one can predict.

But taken together, the research points consistently in the same direction. Here's how we arrive at the 20% estimate:

The option value at origination represents the expected value of the refinancing right across all possible future rate paths, some where rates drop and the borrower benefits, some where rates stay flat or rise and the option goes unused. It's not a guarantee of savings; it's the average value of having the right to act when conditions are favorable. On a typical mortgage ($400K-$600K), this is 4-4.6% of the balance at origination, roughly $18,000-$27,000.

- **The complete non-actors (Keys, 2016):** 20% of qualified borrowers never refinance when they should, losing the full option value. That alone accounts for 20% of the total option value across all borrowers.
- **The partial losses (Agarwal/Rosen/Yao, 2016):** 57% of borrowers who DO refinance do it suboptimally, acting late, at the wrong rate, or both. These partial losses add to the total.
- **The ongoing overpayment (Berger, 2024):** The slowest borrowers overpay by 60bp/year, $3,000/year on a $500K mortgage.

Combining complete losses with partial losses, **a conservative estimate is that the average borrower loses roughly 20% of their refinancing option value to inattention**. On a $500,000 mortgage with 4.45% option value ($22,244), that's approximately **$4,400-$5,000**. The real number is likely higher, since the 20% floor only counts the complete non-actors and doesn't fully account for the partial losses from the majority who act late.

This is a conservative estimate, and the attentiveness data is based on the general population of borrowers. LOs with strong client relationships likely have more attentive clients than average. We don't have precise attentiveness data for broker clients specifically, so the 20% figure is meant as a general baseline.

The point is not to promise clients a precise dollar figure. It's to show the scale of what's at stake. A client enrolled in a monitoring system from purchase, with near-continuous attentiveness and a research-backed threshold, is positioned to capture the full option value across whatever rate environment unfolds. Either way, they're better off with someone watching than without.

Putting it together: what matters more, the rule or the attention?

The ADL paper (Tables 1, 5-6) compares the optimal refinancing threshold against two simpler rules: the second-order approximation (a square-root rule that's close to optimal but doesn't require the full model) and the PV rule (the simple break-even that most financial advisors recommend). Here's what the data shows:

Mortgage Size	Exact Optimal	Second-Order Rule	PV Rule	Loss from Second-Order	Loss from PV Rule
$100,000	193 bps	163 bps	76 bps	$80	$2,897
$250,000	139 bps	123 bps	44 bps	$92	$9,859
$500,000	118 bps	106 bps	33 bps	$123	$22,244
$1,000,000	107 bps	97 bps	27 bps	$189	$47,531

Source: ADL (2013), Tables 1, 5, and 6, pp. 603 and 609

Look at the $500,000 row. The second-order rule is 12 basis points off from optimal (106 vs 118), and the total loss from that difference is just **$123**. Meanwhile, the PV rule is 85 basis points off (33 vs 118), and the loss is **$22,244**. The closer your rule is to optimal, the smaller the loss, and it drops off fast.

Common LO break-even targets of 24-36 months correspond to roughly 84-128 basis points (on a $400K mortgage at 7% with $8K in costs). That puts them in the same range as the second-order rule, far closer to optimal than the PV rule. So the marginal loss from using a 24-36 month rule instead of the full ADL/Berger model is likely in the hundreds of dollars, not thousands.

The conclusion: The refinancing rule matters, but for LOs already using a 24-36 month target, the improvement from switching to the ADL/Berger model is modest. BPS Hawk's suggestion is more

precise because it accounts for factors the simple rule misses (option value, tax asymmetry, moving probability, rate volatility), but it's a refinement of a reasonable baseline, not a correction of a catastrophic error. Note that the model's suggestion can go higher than 36 months depending on the client's situation, sometimes into the 45-48 month range for clients with smaller balances, higher closing costs, or shorter expected ownership. That's normal. The suggestion is personalized, and a higher number just means the math says to wait for a bigger rate drop before acting.

The catastrophic error is not paying attention. The research consistently shows that inattention, not the wrong formula, is what drives the bulk of losses. And the uncertainty of future rate movements can't be calculated away. No one knows where rates are going. What can be controlled is whether someone is watching and ready to act when the math says it's time. With BPS Hawk, the LO has a research-backed rule and near-continuous attentiveness, and that combination is what captures the value that most borrowers leave on the table.

The $250/month example: why being a year late costs more than you think

Consider a $400,000 mortgage. The refinancing option value at origination is approximately $17,800 (4.45% of balance). Say the threshold is hit and your client would save $250/month by refinancing. They check rates once a year, which is actually far more attentive than the average borrower (who checks every 3-4 years).

Just the direct cost of being 12 months late: $250 x 12 = **$3,000**. That's already 17% of the total option value on this loan, nearly the entire 20% estimate, and this is for a borrower who is more attentive than average.

But it's worse than that. Rates don't wait. The window that triggered the opportunity may have only been open for a few weeks or months before bouncing back up. By the time your client checks a year later, that rate is likely gone. They didn't just lose 12 months of savings. They may have lost the entire opportunity. If rates don't come back down for 2-3 years (common), that's $250 x 24-36 months = $6,000-$9,000. If they never come back, your client loses the full lifetime value, potentially the entire $17,800.

The math is simple: $250/month x 12 months late = $3,000 assumes the rate is still there when they finally check. It often isn't. The real cost of checking once a year is the probability of missing the window entirely, multiplied by everything they would have saved over the life of the loan.

---

3. Why This Matters Most for Purchase Clients

You might think this is only about refinancing. It's not. Every purchase client you close is a monitoring client from day one.

For the client:

The moment they close on a $500,000 purchase, you can tell them: "Your mortgage has about 4-4.6% of its value, over $22,000, in built-in refinancing option value. That exists from day one.

Research shows the average person loses at least 20% of that just from not paying attention. I'm going to make sure that doesn't happen to you."

That's a specific, research-backed number, not a vague promise. It changes the conversation from "I'll keep an eye on things" to "here's exactly what's at stake, and here's exactly what I'm doing about it."

How BPS Hawk changes the refinancing conversation

The biggest problem in refinancing isn't the math. It's the psychology. Borrowers are constantly worried about whether rates will go up or down. They hesitate, second-guess, wait for a better deal, and miss the window. The research is clear: you should not have a directional bias on rates. Trying to time the market (waiting because you think rates might drop further, or rushing because you think they might rise) is what causes losses and indecision (ADL, 2013; Berger, 2025).

BPS Hawk solves this by making the decision ahead of time:

1. **You set the target at purchase** (or whenever the client is enrolled). The system calculates a personalized break-even suggestion based on the client's actual loan. This is the number, the threshold where the research says refinancing makes financial sense.
1. **The client is already bought in.** When you explain the target at enrollment ("if rates drop to where your break-even is 24 months or less, I'm going to call you"), the decision framework is established before the opportunity even arises. There's no deliberation when the moment comes. The client trusts the number because you explained it upfront.
1. **You get near-perfect attentiveness.** BPS Hawk monitors rates daily. When a client's rates approach their threshold, you get alerted. You update the numbers, verify the opportunity is real, and call the client. The research shows the average borrower evaluates refinancing once every 3-4 years. With BPS Hawk, you're evaluating continuously. That's the difference between capturing 80%+ of the option value and losing most of it.
1. **No more guessing.** The rate drops, you check the client's rate table, the math is already done. You're not calling clients with a hunch. You're calling with a specific number: "Your break-even is 22 months. That's below your target. Here are the rate options." The conversation is about execution, not speculation.

A note on inbound refinance inquiries: This system is designed for proactive monitoring, setting the threshold before the opportunity arises, not reacting after a client calls asking about rates they saw in the news. For inbound refinance inquiries where the client is already motivated, the rate table still works as a tool, but the pitch is different. If you need help framing those conversations, reach out to support and we're happy to help.

How this applies to you year-round

- **Year-round application.** Even when rates are flat and nobody is refinancing, you're still originating purchases. Every purchase is a new client enrolled in BPS Hawk.

- **Competitive advantage at origination.** No other broker is telling purchase clients the dollar value of their refinancing option, approximately 4-4.6% of their balance. You are. That's a concrete number they've never heard before.
- **Client retention.** A client enrolled in BPS Hawk has a reason to stay with you. If they go to another broker who doesn't monitor, they lose access to that value. You become the broker they can't afford to leave.
- **Learning effect.** Borrowers who have been through the refinancing process once make better decisions the second time, smaller errors, less delay (Agarwal, Rosen & Yao, 2016). Your first refinance with a client builds trust and makes every subsequent one easier.

---

4. How BPS Hawk's Suggestion Works: The Research Behind It

> ***Note:*** *The B/E suggestion is currently in beta. It is based on published, peer-reviewed academic models and is provided for informational purposes only. It is not financial advice. LOs should use their own judgment when acting on the suggestion.*

BPS Hawk computes a personalized break-even suggestion for each client based on their actual loan parameters: mortgage size, holding period, closing costs, tax rate, and more. This section explains the research the suggestion is built on and how the two models work together.

Your Current Rule Is Probably Close

Most LOs use a simple break-even rule: divide closing costs by monthly savings and see how many months to recoup. Common targets are 24-36 months. **If that's your rule, keep using it — the research confirms it's close to optimal for most loan sizes.** BPS Hawk's suggestion is a refinement, not a correction of a catastrophic error. The marginal improvement from switching to the full model is typically a few hundred dollars on a $500K mortgage. What matters far more is that you're attentive. The suggestion is just a more precise starting point.

That said, the simple rule is still suboptimal because it doesn't account for several important factors.

To see why, consider a concrete example: $400,000 mortgage at 7%, $8,000 in closing costs (2%), 30-year term.

Break-Even Target	Monthly Savings Needed	Rate Drop Required
24 months	$333/mo	128 basis points
30 months	$267/mo	101 basis points
36 months	$222/mo	84 basis points
ADL optimal	$327/mo	~125 basis points

The ADL optimal threshold for this mortgage is approximately 125 basis points, which translates to a 24.5-month simple break-even. So if you're already using a 24-month rule, you're accidentally close

to optimal for this loan size. A 36-month rule is more lenient (triggers at smaller rate drops), which means refinancing slightly too early relative to optimal.

Why the simple rule is still suboptimal: It doesn't account for the option value of waiting (the possibility that rates might drop further after you refinance), tax asymmetry (mortgage interest is deductible but closing costs are not), the probability of moving, interest rate volatility, or the time value of money. The ADL model accounts for all of these. For some clients, especially those with smaller balances, higher closing costs, or shorter expected ownership, the difference between the simple rule and the optimal threshold is significant.

The cost of using a simple break-even rule instead of the optimal one:

Mortgage Size	Cost of Wrong Rule	% of Balance
$100,000	$2,897	2.90%
$250,000	$9,859	3.94%
$500,000	$22,244	4.45%
$1,000,000	$47,531	4.75%

Source: ADL (2013), Tables 5-6, p. 609. Chen & Ling (1989) independently showed that simple break-even models understate the required rate differential by 37-311 basis points depending on conditions.

These loss numbers assume perfect attentiveness; the borrower watches rates continuously and acts the moment their rule triggers. As Section 2 shows, that's not how it works in reality. The rule matters, but the attention problem matters more.

The ADL Model (Agarwal, Driscoll & Laibson, 2013)

Published in the *Journal of Money, Credit and Banking*, the ADL paper solves the question: "How far do rates need to drop before it makes sense to refinance, accounting for the possibility that rates might drop even further?"

This is what mathematicians call an "optimal stopping problem." The borrower has to decide: act now, or wait? Acting has a known cost (closing costs) and a known benefit (lower rate). Waiting has an unknown outcome; rates might go lower (better) or bounce back (worse). The ADL model finds the exact point where acting beats waiting, given the uncertainty.

What it takes as inputs (all specific to the client):

- Mortgage balance
- Current rate
- Remaining term and original term
- Closing costs (fixed costs + points)
- Marginal tax rate (mortgage interest is deductible, closing costs are not; this asymmetry matters)
- Expected ownership time (probability of moving)

- Interest rate volatility (how much rates bounce around)
- Discount rate and inflation

What it produces:

A single number: the optimal rate reduction threshold in basis points. When rates drop by at least this much below the client's current rate, the math says refinancing makes sense.

The optimal threshold varies by client:

Mortgage Size	Optimal Threshold
$100,000	193 basis points
$250,000	139 basis points
$500,000	118 basis points
$1,000,000	107 basis points

Source: ADL (2013), Table 1, p. 603

A broker telling every client "wait for a 1% drop" is wrong for most of them. The right threshold depends on the client's specific loan.

The Berger Extension (Berger, Milbradt, Tourre & Vavra, 2025)

Published in *AER: Insights*, the Berger paper builds on ADL by adding one critical real-world factor: borrowers don't check rates continuously. They check sporadically, on average once every 3-4 years, based on data from 780,000 Fannie Mae mortgages and 250,000 Equifax borrowers.

The key insight: If a borrower rarely gets to check rates, then when they do check, they should be more willing to act, even on a smaller rate drop. Why? Because they don't know when their next chance will be. The option value of waiting is lower when you can't monitor continuously, because "waiting" might mean years of not looking.

Berger quantified this precisely. The ADL threshold assumes perfect attention; the borrower can act at any moment. The Berger-adjusted threshold is lower, reflecting realistic inattention. In their calibration, inattention reduces the threshold by 54-65% compared to ADL.

Example with ADL baseline parameters ($250,000 mortgage):

- ADL threshold (perfect attention): ~146 basis points
- Berger threshold (checking once every 4.3 years): ~61 basis points
- The inattention adjustment cuts the threshold by more than half

What Berger adds as an input:

One new parameter beyond what ADL uses: the attention rate. This is how often the borrower (or their broker) evaluates whether to refinance. More attention = higher threshold (you can afford to be pickier). Less attention = lower threshold (take the opportunity when you see it).

How They Work Together in BPS Hawk

1. **ADL computes the optimal threshold** assuming continuous monitoring. This is the theoretical best case. This uses the client's actual loan parameters: balance, rate, term, costs, tax rate, ownership time, and rate volatility. The ADL calculation is done at "par" (no points, f = 0), which means it computes the threshold based only on fixed closing costs, ignoring the points cost of any specific rate option.
2. **Berger adjusts that threshold** based on how often rates are checked. BPS Hawk defaults the attention rate to 365 times per year — daily — because the system monitors daily. At 365, the Berger-adjusted threshold is essentially equal to the ADL optimal (x*). This means the default suggestion is the research-backed optimal threshold under perfect attention. The suggestion is personalized per client — all inputs including the attention rate are specific to each client record and can be overridden at any time.
3. **The threshold is converted to a break-even in months.** The system takes the Berger-adjusted rate threshold and computes: at that target rate, what are the monthly payment savings? Then: total closing costs / monthly savings = break-even in months. This is the suggestion shown on the client detail page.
4. **That one number is applied across the entire rate table.** The suggestion is computed once at par and used as the target for every row, regardless of how many points that row costs. Each row on the rate table has its own break-even (which includes that row's points in the total cost), and any row where the break-even meets or beats the par-based suggestion is highlighted. This is the simplification that makes the system work: one target number, consistent across all rate options. The LO doesn't need to compute a different threshold for each points level.

If you want a lower threshold (smaller rate drop required to trigger action): Lower the attention rate on the client record. A lower attention rate tells the model the borrower checks less frequently, which pushes the threshold down — the model says act sooner when you check infrequently. The sensitivity table shows this in real time: as you move down the rows (lower attention), the threshold and break-even suggestion decrease. Importantly, the W*M column shows the option value barely changes across rows — there is no materially wrong answer within a reasonable range. Use your judgment for each client.

A note on how the break-even scales down the rate table: Technically, each rate option on the table represents a different "deal": different rate, different points cost, different closing costs. If you ran the full ADL model separately for each row (using that row's V values), you'd get slightly different optimal answers per row. The V method tends to favor going further down the table (lower rates with more points) because it accounts for the full economic picture including how the option value changes at each rate level. The simple break-even has a slight bias toward staying closer to the top of the table (fewer points, higher rate).

I've calculated the maximum loss from this bias at approximately 0.125% of the loan amount, about $500 on a $400,000 mortgage. And the bias only shows up meaningfully when you're looking at very expensive options with roughly 1.5%+ in points, which are options that most borrowers don't

want to take anyway. For the rate options that borrowers actually consider, the simple break-even and the V method give the same answer. I can demonstrate this with the data upon request.

The primary reason for this bias is the tax deductibility of points. The ADL model (Step 2) reduces the effective cost of points because they are partially tax-deductible over the loan term. The simple break-even uses points at face value. So when you go further down the table to a row with 1.5% in points, the simple BE treats the full points cost as the expense, while the model sees a lower after-tax effective cost. This is why the V method says "going further down is cheaper than you think" because it accounts for the tax benefit, while the simple BE doesn't. Additionally, the option value of the new loan is slightly lower at a lower rate (less room for rates to drop further), which also plays into the V calculation favoring those rows, though this effect is small compared to the tax adjustment on points.

Why we show pre-tax numbers on the table: The table shows actual payment amounts and actual closing costs, what the borrower will see on their statement and at the closing table. This is simpler to understand and is what LOs and clients are used to working with. The tax adjustments happen inside the ADL model when computing the suggestion threshold, so the decision is tax-aware even though the display is pre-tax. Showing after-tax numbers would be more precise but harder to explain and verify, and since the suggestion already accounts for taxes, the display doesn't need to.

The takeaway: using one target number across the table costs at most ~$500 in the worst case, while the attention problem costs $4,400+. The simplification is well worth it for usability.

Why the suggestion is in months, not basis points:

LOs already think in break-even months. It's the standard language of the industry. A client understands "24 months to recoup your costs" far more easily than "128 basis points." It also gives the LO one consistent number to use as the target across the rate table: any rate option where the break-even meets or beats the suggestion is highlighted. This keeps the conversation simple.

Using a single break-even number in months is a simplification. The ADL/Berger model works in basis points and accounts for discounting, tax adjustments, and option value dynamics. The simple break-even (total costs / monthly savings) doesn't capture all of those factors perfectly, so it doesn't scale identically across every loan scenario. But in my estimation it scales very well. The difference is typically a month or two at most. And as Section 2 shows, the refinancing rule matters far less than attentiveness, so a small imprecision in how the rule is presented is not something to worry about. The whole point of expressing the suggestion in months is to make the system usable: one number, consistent across the table, easy to explain to clients. If you want a deeper explanation of how the conversion works, reach out to support and I'm happy to walk through it.

A note on taxes: The break-even numbers on the rate table (total costs, monthly payment savings, and the simple break-even in months) are all computed before tax. The monthly payment shown is the actual mortgage payment, not the after-tax cost. However, the ADL model that generates the suggestion threshold does account for taxes internally (the tax deductibility of mortgage interest and points is built into the formula; see Appendix: ADL Formula, Steps 2-3). So the suggestion itself is tax-aware, even though the table displays pre-tax numbers. This is a deliberate choice: LOs and

clients think in terms of actual payments, not after-tax payments, so the table shows what they're used to seeing.

What the Break-Even Suggestion Table Shows

On each client's detail page (refinance and purchase), the system displays a Break-Even Suggestion table. This table lets the LO see how the suggested break-even changes under different assumptions and inputs:

Rows based on attention rate: The table shows the suggested break-even at several attention levels. This lets the LO see how the suggestion shifts depending on how attentive the monitoring is. Higher attention = higher threshold (you can be pickier) = longer break-even in months. Lower attention = lower threshold (act when you can) = shorter break-even. The client's actual attention rate (defaulted to 365/daily) is highlighted.

The x* row (perfect attention): The top row shows the ADL optimal threshold, the theoretical best assuming continuous monitoring. This is the upper bound.

Custom BPS input: The LO can enter any basis point level manually and see the break-even that corresponds to it. This lets them test different scenarios or use their own threshold if they prefer.

The W*M column (option value): Each row also shows W*M, a dollar value computed using the ADL welfare formula. This is the value of the refinancing option at origination**, meaning at the moment the loan is taken out, when the market rate equals the client's rate and there is no rate gap yet. It represents: "if I follow this particular rule from day one, what is the expected value of having the right to refinance?"

You'll notice W*M is similar across rows even when the thresholds vary significantly. This seems counterintuitive — shouldn't waiting for a bigger drop capture more value? The reason it doesn't is that a higher threshold is a bet: rates need to drop far enough to hit it. A lower threshold acts more often on smaller drops. A higher threshold acts less often but saves more each time — if rates get there. At origination, these two effects roughly cancel out. You can cash in the option sooner for a smaller gain, or hold it for a bigger gain that may or may not arrive. The expected value is similar either way. This is why the rule matters less than whether you're watching.

At x *(perfect attention), W*M is highest. That's the theoretical maximum. At lower thresholds, W*M is slightly lower. The difference is modest, which is precisely why the rule matters less than attentiveness. Using the default attention rate of 6 times per year instead of perfect attention results in about a 3% reduction in option value, a very small tradeoff for a more realistic suggestion.

Important note: the W*M values assume perfect attentiveness at whatever threshold is shown. The ADL formula computes the value of following a given rule while watching rates continuously. We don't currently have a way to compute the option value under imperfect attentiveness. The Berger extension adjusts the threshold for inattention, but the welfare calculation itself still assumes the borrower acts every time that threshold is hit.

Works the same for purchase clients: The table works identically for purchase clients, using the par rate from their purchase rate sheet as the starting point instead of their current mortgage rate. This lets the LO show a purchase client the option value built into their new mortgage from day one.

Attention vs. Rule: Which Matters More?

The ADL welfare table shows that different rules produce similar option values at origination. The W*M column on the sensitivity table makes this visible: even as the threshold changes significantly between attention levels, the dollar value of the option stays roughly the same. The rule shifts the threshold, but the expected value of following any reasonable rule is similar, because a lower threshold means acting more often on smaller drops, which roughly compensates for the smaller per-action savings.

Where the real money is lost is not in which rule you follow. It's in whether you follow any rule at all. The ADL table shows that using the wrong rule costs $2,900-$47,500, but those numbers are for the PV rule (a 65-70 month break-even that no LO actually uses). For LOs using a 24-36 month target, the marginal loss from the rule itself is in the hundreds of dollars (see the second-order rule comparison in Section 2). Meanwhile, 20% of borrowers lose the entire option value by never acting, and 57% of those who do act get the timing wrong.

A borrower using a simple break-even rule but checking monthly will outperform a borrower using the perfect formula who checks once every 4 years.

This is exactly why BPS Hawk focuses on monitoring first. The rule matters, and ours is better, but fixing the attention problem is where the real money is.

---

5. What to Tell Your Clients

For detailed scripts and walk-throughs on how to present this to clients, see the BPS Hawk Tutorial Guide (.claude/docs/tutorial/).

The Purchase Enrollment Pitch:

"Your new $500,000 mortgage has over $22,000 in refinancing option value built into it. That means if rates drop at the right time and you act at the right moment, you could save that much over the life of your loan. Research shows the average person loses at least $4,000 of that just from not paying attention. I'm enrolling you in our monitoring system right now. When rates hit your personalized threshold, I'll call you. You don't have to think about it."

The Retention Pitch:

"Research across over a million mortgages shows that borrowers without monitoring lose thousands of dollars because they miss refinancing opportunities. As long as you're my client, I'm watching your rate every day. That's worth real money to you."

---

6. Should My Client Buy Points?

The same framework that answers "when should I refinance?" also answers "should I take points?", for both refinance and purchase clients.

How it works for refinance

Each row on the rate table represents a different rate option, and each option has a different points cost. A lower rate costs more points (higher upfront cost), while a higher rate costs fewer points (or even gives a lender credit). The system computes the break-even for every row, including the points cost in the total closing costs.

The LO doesn't need to do a separate analysis. They look at the rate table and check: does this rate option's break-even meet the suggestion target? If a rate with 1 point has a break-even of 22 months and the target is 24 months, the points are worth it. If the break-even is 30 months and the target is 24, they're not. The simplest approach: find the lowest break-even on the table that meets the target. That's the best option.

Under the hood, each row's tax adjustment is computed directly: Tax Adj = -(Points Cost × Tax Rate). This accounts for the fact that points are partially tax-deductible, which the simple break-even rule doesn't capture. (The ADL model that generates the suggestion runs once at par — the per-row tax adjustment is a simpler formula applied independently to each row.) The result is a more accurate assessment of whether points make sense for that specific client.

Compare mode on refinance: The rate table also has a compare mode. The LO can click any row to set it as the "base" rate, and the table shows the break-even of going from that base rate to each lower rate. This answers: "if I'm already considering the 6.25% option, is it worth paying extra points to get 6.00%?" The compare break-even between those two rows tells you. On refinance, this will give you the same answer as just reading the main break-even column. If the lowest break-even that meets the target is 6.00%, then comparing 6.25% to 6.00% will also show it's worth it. The compare mode is there for LOs who want to see the math between specific rate options.

How it works for purchase

On a purchase, there is no current loan to compare against. The client is choosing between rate options at origination: do they take the par rate, or buy down to a lower rate with points?

The rule is simple: **set the highest available rate as the base.** From there, compare mode shows the break-even of buying down to each lower rate. Whichever option has the lowest break-even that still meets the suggestion target is the best choice. If no option beats the suggestion target, the base rate (highest rate, fewest points) is the right answer — the buydown isn't worth it.

Refinance: using Compare mode to pick the best rate

The same compare logic applies to refinance when you're evaluating which rate option to take. Set the highest available rate as the base and compare down from there. The one exception: if a rate option has a negative total closing cost (the lender credit exceeds all costs — effectively free

money), you would not use that as your comparison base, since it distorts the break-even math. In that case, set the highest rate without a negative total cost as your base.

Beyond that, the same rule applies: whichever option has the lowest break-even that meets the suggestion target is best. If none do, the base rate is the right choice.

The Tax Adj column

The rate table includes a **Tax Adj** column that adjusts the break-even for the tax benefit of paying discount points. When the client's marginal tax rate is on file, the system reduces the net cost by the expected deduction — shown as a negative number (savings). For purchase clients, points are fully deductible in Year 1. For refinance, the deduction is technically amortized over the loan term, but the system applies the simplified full-deduction as an approximation. The **Adj B/E** column reflects this adjustment and is the primary number used for comparing rate options. See the Appendix for the full formula and deductibility details.

The bottom line on points

Set the highest rate as the base. Compare down. Lowest **Adj B/E** that meets the target wins (or Spl B/E if no tax rate is set). If nothing meets the target, the base rate is the answer. This works the same way for purchase and refinance.

---

7. How BPS Hawk Gets Accurate Rates

We source the best available base pricing for both Conventional and FHA loans. Each rate level has a corresponding base points value. The system then applies the client's full LLPA profile on top: credit score, LTV, property type, occupancy, and state adjustments are all stored and applied automatically to produce the actual rate options for that specific client.

If the rates feel slightly off for your market, they will be off in the same direction consistently — meaning the error is systematic, not random. To account for this, we can apply a tuning adjustment to your account to calibrate the grid to your pricing. Based on our research, rates at each level don't vary significantly between lenders on a given day, so once tuned, the system produces very accurate results.

---

8. The Numbers at a Glance

Stat	Source	What It Means
Option value: 4-4.6% of balance	ADL (2013), Tables 5-6	What's at stake on every mortgage
At least 20% of option value lost to inattention	Keys (2016), Berger (2024), ARY (2016)	~$4,400+ on a $500K mortgage

20% of qualified borrowers never refi	Keys, Pope & Pope (2016), 376K mortgages	1 in 5 of your clients is losing everything
Median PV loss for non-refinancers: $11,568	Keys, Pope & Pope (2016)	On ~$200K avg mortgages
84% ignored a free pre-approved offer	Keys, Pope & Pope (2016), nonprofit experiment	Pure inattention, not inability
57% of refinancers do it suboptimally	Agarwal, Rosen & Yao (2016), 271K refis	Even those who act usually get it wrong
Avg refinancing error: 60 basis points	Agarwal, Rosen & Yao (2016)	Refinancing at a rate higher than optimal
2.66 months of missed opportunity before acting	Agarwal, Rosen & Yao (2016)	The window sits there while they wait
Borrowers evaluate refi once every 3-4 years	Berger (2025), Berger (2024)	The attention problem is massive
45% of borrowers almost never evaluate	Berger et al. (2024), 250K borrowers	Nearly half aren't watching at all
Slowest borrowers overpay 60bp/year	Berger et al. (2024)	$3,000/yr on a $500K mortgage
Cross-subsidies are regressive	Berger et al. (2024)	Less wealthy clients subsidize more wealthy ones
$5.4 billion forgone savings/year nationwide	Keys, Pope & Pope (2016)	The national scale of the problem
Brokers don't improve outcomes without a system	Agarwal, Rosen & Yao (2016)	Without monitoring, you're guessing too
Avg points taker loses ~$700	Agarwal, Ben-David & Yao (2017), 312K mortgages	Better refi behavior = skip points
Points negative NPV in 98-100% of months	Agarwal, Ben-David & Yao (2017)	Almost never the right call
2nd-time refinancers make smaller errors	Agarwal, Rosen & Yao (2016)	Learning effect, first refi builds trust
Missing refi windows increases default risk	Chang & Yavas (2009), 3,785 loans	Monitoring prevents loss, not just saves money

---

Academic Sources

Agarwal, S., Driscoll, J.C., & Laibson, D.I. (2013). "Optimal Mortgage Refinancing: A Closed-Form Solution." *Journal of Money, Credit and Banking*, 45(4), 591-622.

Berger, D., Milbradt, K., Tourre, F., & Vavra, J. (2025). "Optimal Mortgage Refinancing with Inattention." *AER: Insights*, 7(4), 497-516.

Berger, D., Milbradt, K., Tourre, F., & Vavra, J. (2024). "Refinancing Frictions, Mortgage Pricing and Redistribution." NBER Working Paper No. 32022.

Keys, B.J., Pope, D.G., & Pope, J.C. (2016). "Failure to Refinance." *Journal of Financial Economics*, 122(3), 482-499.

Agarwal, S., Rosen, R.J., & Yao, V. (2016). "Why Do Borrowers Make Mortgage Refinancing Mistakes?" *Management Science*, 62(12), 3494-3509.

Agarwal, S., Ben-David, I., & Yao, V. (2017). "Systematic Mistakes in the Mortgage Market and Lack of Financial Sophistication." *Journal of Financial Economics*, 123(1), 42-58.

Chang, Y. & Yavas, A. (2009). "Do Borrowers Make Rational Choices on Points and Refinancing?" *Real Estate Economics*, 37(4), 635-658.

Chen, A.H. & Ling, D.C. (1989). "Optimal Mortgage Refinancing with Stochastic Interest Rates." *AREUEA Journal*, 17(3), 278-299.

Appendix

A note on the Appendix: The formulas and math below are provided for reference and transparency. They are not required reading. The system handles all of this automatically. The LO only needs to understand the suggestion number and the rate table. The Appendix is there for anyone who wants to understand what's under the hood or verify how the numbers are produced.

Model Assumptions

The ADL/Berger model that powers BPS Hawk's suggestion relies on several assumptions. These are standard in the academic literature and produce results that closely match more complex numerical models (within 10 basis points per Chen & Ling, 1989), but LOs should be aware of them:

1. **Interest rates follow a random walk (Brownian motion).** The model assumes rates move randomly with no predictable trend. They're as likely to go up as down at any given moment. In reality, rates may have some tendency to revert toward a long-run average (mean reversion). This simplification is necessary to get a closed-form solution, and the ADL paper notes it closely matches numerical methods that don't make this assumption.
1. **Constant parameters.** The discount rate, volatility, moving probability, and tax rate are assumed to stay the same over the life of the loan. In reality, these change. The model captures a snapshot based on current conditions.

1. **Risk neutrality.** The borrower is assumed to care only about expected dollar values, not about risk. In reality, some borrowers are more risk-averse and may prefer to lock in savings sooner rather than gamble on rates dropping further.
1. **Infinite refinancing option.** After refinancing, the borrower gets a new mortgage with the same option to refinance again. This is approximately true for U.S. fixed-rate mortgages.
1. **Tax rate is constant for the life of the loan.** The marginal tax rate (used to compute the after-tax cost of closing and the tax benefit of mortgage interest deduction) is assumed to remain the same throughout. In reality, a borrower's tax bracket can change due to income changes, tax law changes, or life events. The model uses the current marginal tax rate as the best available estimate (see note on tax rate below).
1. **Attentiveness is constant.** The Berger extension assumes the borrower checks rates at a fixed frequency (Poisson process). With BPS Hawk monitoring daily, this assumption is conservative. The system effectively provides near-continuous attention.

How the suggestion is calculated (step by step)

For each client, BPS Hawk computes the break-even suggestion as follows:

1. **Gather client inputs.** From the client record: mortgage balance (M), current rate (i0), remaining term, original term, marginal tax rate (tau), expected ownership time (mu). From market defaults: discount rate (rho), inflation (pi), interest rate volatility (sigma).
1. **Compute fixed closing costs.** FCC = fixed closing costs percentage × mortgage balance (default 2%). This is the "F" input to the model.
1. **Run the ADL model (7 steps).** Using the inputs above with points = 0 (par rate), the model computes the optimal refinancing threshold x* in basis points. This accounts for the option value of waiting, tax asymmetry, moving probability, and rate volatility.
1. **Run the Berger extension (4 additional steps).** Using the ADL results plus the client's attention rate (default 6 times/year), the model adjusts the threshold downward for realistic attentiveness. This produces theta_B, the inattention-adjusted threshold in basis points.
1. **Convert to break-even months.** The system computes: at a rate of (current rate − theta_B), what is the monthly payment? Monthly savings = current payment − new payment. Break-even = total closing costs / monthly savings. This is the suggestion shown on the client detail page.

Why this specific conversion? The ADL/Berger model produces theta_B in basis points. Converting to months via simple break-even (FCC / monthly savings at the trigger rate) is the natural way to express that threshold in LO-friendly terms. It asks: "if the market rate dropped by exactly theta_B right now, how many months would it take to recover closing costs at that savings rate?" That number is the floor. Any rate option on the table with a break-even at or below that floor is meeting or beating the ADL/Berger threshold. The complexity happens in the model. The output is expressed simply.

1. **Apply to the rate table.** Any rate option on the table where the break-even (including MI) meets or beats the suggestion is highlighted. The LO sees at a glance which rates are worth pursuing.

The LO can override the suggestion with a manual target if they prefer a different threshold for a specific client.

For purchase clients: The suggestion calculation is identical, with three input differences:

- **i0 = par rate from the purchase pricing grid.** There is no current mortgage rate. The par rate (highest available rate, zero points) is used as the baseline.
- **Gamma = full loan term.** The loan hasn't started, so there is no remaining term. A 30-year loan uses Gamma = 30.
- **Monthly savings = payment at par − payment at (par − theta_B), on the same balance M.** This is the savings from buying down from the par rate to the trigger rate. The balance is the same in both calculations because purchase points are paid out of pocket, not financed.

Everything else — the ADL steps, the Berger extension, the option value formula — is the same. Theta_B means: "rates need to drop this far below the purchase rate before refinancing makes sense." MI Diff is always zero when computing the suggestion for a purchase client, because there is no existing loan to compare against.

ADL Formula (Steps 1-7)

The ADL model computes the optimal refinancing threshold x* assuming perfect attentiveness. All inputs are specific to the client.

Inputs:

- M = mortgage balance ($)
- F = fixed closing costs ($)
- f = points (decimal, e.g., 0.01 = 1 point). At par, f = 0.
- tau = marginal tax rate (decimal, e.g., 0.28)
- mu = expected ownership time (decimal, from client record)
- i0 = current mortgage rate (decimal, e.g., 0.07). For purchase clients, this is the par rate from the pricing grid.
- Gamma = years remaining on current mortgage. For purchase clients, this is the full loan term.
- N = original term in years (or new loan term for purchase)
- sigma = interest rate volatility (market-wide default)
- rho = discount rate (market-wide default)
- pi = inflation rate (market-wide default)

Why the suggestion is computed at par (f = 0): The suggestion is meant to be one target number that applies consistently across every row on the rate table. If you ran the ADL model separately for

each row using that row's actual points cost (f = 0.5%, 1%, 1.5%, etc.), you'd get a different threshold per row — which defeats the purpose of having a single target. Running at par produces a neutral baseline: "ignoring which specific rate option the client picks, what rate drop does the math say is worth acting on?" Each row then competes against that baseline via its own Adj B/E calculation. The par threshold is the benchmark. The rows are the competitors.

Steps:

```
Step 1: theta = mu + 0.10

Step 2: kappa = F + f * M * (1 - tau/(theta+rho+pi) *
((1-exp(-(theta+rho+pi)*N))/N * (rho+pi)/(theta+rho+pi) + theta))
         At par (f=0): kappa = F (the points term vanishes)

Step 3: C = kappa / (1 - tau)

Step 4: lambda = mu + i0 / (exp(i0 * Gamma) - 1) + pi

Step 5: psi = sqrt(2 * (rho + lambda)) / sigma

Step 6: phi = 1 + psi * (rho + lambda) * C / M

Step 7: x* = (-1/psi) * (phi + W(-exp(-phi)))
         where W is the Lambert W function (principal branch)
         x*_bps = x* * 10000
```

x *is negative. It represents the rate drop needed. For example, x* = -0.0125 means rates need to drop 125 basis points below the client's current rate.

What each step does in plain terms:

- Step 1: Adjusts for the probability of a future refinancing event that would let you deduct remaining points
- Step 2: Computes the true after-tax cost of refinancing, accounting for the tax deductibility of points over time
- Step 3: Converts to "interest-equivalent dollars" (because mortgage interest is deductible but closing costs aren't)
- Step 4: Computes how fast the mortgage value declines from moving, principal paydown, and inflation
- Step 5: Computes how fast the refinancing option value decays as rates rise, driven by volatility and discounting
- Step 6: Packages the cost burden into a single dimensionless number
- Step 7: Solves for the optimal threshold using the Lambert W function, the exact point where acting beats waiting

Berger Extension (Steps 8-11)

The Berger extension adjusts the ADL threshold for borrower inattention. It takes all the ADL results plus one new input: alpha (the attention rate, how many times per year the borrower evaluates refinancing).

Steps:

```
Step 8:  beta = sqrt(2 * (rho + lambda + alpha)) / sigma

Step 9:  gamma = (rho + lambda) * (psi + beta) / alpha

Step 10: omega = C / M

Step 11: theta_B = -[(rho+lambda)*omega + 1/(psi+gamma) + (1/psi)*W(z)]
         where z = (-psi/(psi+gamma)) * exp((-psi/(psi+gamma)) * (1 +
(rho+lambda)*(psi+gamma)*omega))
         theta_B_bps = theta_B * 10000
```

theta_B is the inattention-adjusted threshold. It is smaller (less negative) than x*, meaning inattentive borrowers should act on smaller rate drops when they do check.

What each step does:

- Step 8: Computes a new decay rate that includes attention events (above the threshold, attention events trigger refinancing)
- Step 9: Captures the effect of finite attention on the threshold. Goes to zero as attention increases (recovering ADL), goes to infinity as attention drops to zero
- Step 10: Closing cost as a fraction of balance
- Step 11: Solves for the inattention-adjusted threshold using the Lambert W function

Limiting cases:

- As alpha → infinity (perfect attention): theta_B → x* (recovers the ADL threshold)
- As alpha → 0 (never checks): theta_B → the flow cost of refinancing (just the annualized closing costs, no option value, because the borrower can't act on it)

Option Value Formula (W*M)

The W*M column on the sensitivity table shows the expected value of the refinancing option at each threshold:

```
W(x) = (C/M + x/(rho + lambda)) / (1 - exp(-psi * x))

W*M = W(x) * M
```

At x *(perfect attention), W*M is highest. That's the full option value. At lower thresholds (less attentive), W*M is lower because less of the option value is captured.

Note: W*M assumes the borrower follows the given rule with perfect attentiveness at that threshold. It does not account for the borrower sometimes missing the threshold due to inattention. The Berger extension adjusts the threshold itself for inattention, but the welfare calculation assumes the adjusted threshold is then followed perfectly.

The attention rate setting and why we default to 6

The attention rate is a number on each client's record that represents how often the borrower (or their broker) evaluates whether to refinance. It feeds into the Berger extension of the model. Higher attention = higher threshold (the model says you can afford to be pickier, because you'll catch the next opportunity). Lower attention = lower threshold (act when you can, because you might not check again for a while).

You can see the effect on the Break-Even Suggestion table on each client's page. It shows the suggested break-even at different attention levels. The difference between levels is visible but modest — the W*M column shows that option value barely changes across rows, so there is no materially wrong answer within a reasonable range.

Why we default to 365 (daily): BPS Hawk monitors rates every day, so assuming daily attention is the correct assumption. At 365, the Berger-adjusted threshold is essentially the same as the ADL optimal (x*) — the research-backed threshold under perfect attention. This gives LOs the highest, most research-grounded threshold by default.

If you want a lower threshold: Lower the attention rate on the client record. The suggestion recalculates instantly and the sensitivity table shows the effect. Use this when a client has a specific situation where you want to act on a smaller rate drop. The LO can adjust per client at any time.

The LO can adjust the attention rate per client at any time, and the suggestion recalculates.

Parameters the LO doesn't see:

The ADL/Berger model uses several market-wide parameters that are set by the admin and not visible on the client detail page:

- **Sigma (interest rate volatility):** Default 0.70% per year. How much rates bounce around. Higher volatility means the option to wait is more valuable.
- **Rho (discount rate):** Default 5%. How much the borrower values money today vs. the future. Standard assumption in the literature.
- **Pi (inflation rate):** Default 3%. Erodes the real value of the mortgage over time.

These three parameters feed into the ADL formula and affect the suggestion threshold, but they are the same for all clients and change infrequently. The admin can update them if market conditions shift materially, and the update applies to all clients at once.

The fields the LO does see and can edit per client (unless hidden by the admin) include: fixed closing costs, expected ownership time, appreciation rate, marginal tax rate, attention rate, MI rate, and current monthly MI. These are the client-specific inputs that drive the personalized suggestion. They default to values set by the admin at client creation, but the LO can override them for any individual client.

Note on the tax rate assumption

The model uses the client's marginal tax rate in two places: (1) computing the after-tax cost of refinancing (closing costs are not tax-deductible, but mortgage interest is, so $1 of closing costs is

effectively more expensive than $1 of interest), and (2) computing the present value of future interest savings after the tax deduction.

The assumption is that this tax rate stays the same for the entire life of the loan. In reality, a borrower's marginal rate can change. Income goes up or down, tax brackets change, deduction limits shift. The model uses the current rate as the best available estimate. If a borrower's tax situation changes materially (e.g., they move into a much higher or lower bracket), the LO can update the tax rate on the client record and the suggestion will recalculate.

The impact of the tax rate on the threshold is moderate: ADL (2013, Table 2) shows that moving from a 0% to 35% marginal tax rate changes the optimal threshold by about 10-35 basis points depending on loan size. So even a meaningful change in tax bracket shifts the suggestion by a relatively small amount.

Note on tax deductibility of points and MI

The rate table includes a **Tax Adj** column that adjusts the break-even calculation for borrowers who itemize deductions. Here is what's deductible and when it matters.

The formulas

```
Tax Adj  = -(Points Cost × Tax Rate)   when tax rate > 0 AND borrower is
paying points
Tax Adj  = 0                           when tax rate = 0 OR borrower has a
lender credit

Adj. Cost = Total Cost + Tax Adj + MI Diff
            (Tax Adj is negative, so it reduces Adj. Cost)

Adj B/E  = ceil((Total Cost + Tax Adj + MI Diff) / Monthly Payment Savings)
```

Example: A client in the 24% bracket pays $8,000 in points. Tax Adj = -(8,000 × 0.24) = -$1,920. If MI Diff = $0, then Adj. Cost = $8,000 + (-$1,920) + $0 = $6,080. At $200/month in savings, Adj B/E = ceil($6,080 / $200) = 31 months instead of the pre-tax 40 months.

When the tax rate is zero, Tax Adj = $0, Adj. Cost = Total Cost + MI Diff, and Adj B/E equals what the system previously showed as "BE+MI." There is no behavior change for clients without a tax rate set.

Discount points

Points paid by the borrower are tax-deductible. For a purchase loan, they are fully deductible in Year 1. For a refinance, the IRS requires them to be amortized over the life of the loan. The system applies the full-deduction approximation for both — a slight overstatement for refi borrowers — which is a common simplification when estimating the net cost of points at the time of closing. LOs should encourage clients to confirm the exact treatment with a CPA.

When a client receives a lender credit instead of paying points, there is nothing to deduct. The Tax Adj column shows zero in that case.

Mortgage insurance

MI premiums are **not deductible for the 2025 tax year**. Under the One Big Beautiful Bill Act (signed July 4, 2025), MI deductibility becomes permanent starting with the **2026 tax year**. The system accounts for after-tax MI cost in the Adj B/E calculation — the present value of MI payments is reduced by the client's marginal tax rate. If MI applies (initial LTV > 80%), the tax benefit is built in when a tax rate is on file.

MI deductibility has changed multiple times. Confirm current status with a CPA, particularly for higher-income borrowers where phase-out rules may apply.

Who should have a tax rate on file

Enter the client's combined federal + state marginal tax rate if they itemize deductions and will be paying discount points. If the client takes the standard deduction, is receiving a lender credit, or you're not sure — leave it at zero. The break-even numbers are correct either way; the Tax Adj column just does nothing when tau = 0.

A note on pre-tax vs. after-tax numbers

The rate table has always shown pre-tax monthly payments (the actual payment the client makes), and that does not change. The Tax Adj column adds one additional layer of adjustment — the tax deduction on points — specifically because it is a lump-sum benefit at closing that meaningfully affects the break-even. The monthly payment shown is still the gross payment. This is the right approach: the deduction on points is finite and calculable upfront, while the ongoing tax benefit of mortgage interest deduction is already handled inside the ADL model that generates the suggestion threshold.

Always recommend a CPA conversation

Tax rules change. The Tax Adj calculation is an estimate designed to improve the accuracy of rate comparisons — it is not tax advice. For any client where the tax adjustment changes which rate looks best, have them confirm the deductibility with their CPA before closing.

Mortgage Insurance (MI) Logic

MI affects the break-even calculation. When a borrower's LTV is above 80%, they pay mortgage insurance, which adds to the monthly cost. If refinancing changes the MI situation (e.g., the new loan also has MI, or MI drops off sooner due to appreciation), this affects whether refinancing is worthwhile.

When MI applies: MI is required when initial LTV > 80% (mortgage balance / property value). MI drops to $0 when LTV reaches 80% as the balance amortizes and the home value appreciates.

How MI is calculated:

```
Step 1: Determine monthly MI payment
  - Refi loan:    monthlyMI = mi_rate * M / 12
  - Existing loan: monthlyMI = current_monthly_mi (from client record)

Step 2: Find month T where MI drops off
  For each month t = 1, 2, 3, ...:
    balance_t = M * (1+r)^t - payment * ((1+r)^t - 1) / r
```

```
    value_t   = property_value * (1 + g/12)^t
    If balance_t / value_t <= 0.80: T = t, stop

Step 3: Sum discounted MI payments from month 1 to T
  PV_MI = Sum(t=1 to T) [ monthlyMI * (1 + rho/12)^(-t) * (1 - mu/12)^t ]

  where:
    (1 + rho/12)^(-t) = discount factor (time value of money)
    (1 - mu/12)^t     = survival probability (chance borrower hasn't
moved/prepaid)

Step 4: After-tax adjustment
  PV_MI_aftertax = PV_MI * (1 - tau)
```

How MI enters the break-even:

The system computes the PV of MI for both the existing loan and the refi loan, then takes the net difference:

```
MI_diff = PV_MI_refi_aftertax - PV_MI_existing_aftertax
```

This MI difference is added to the total cost of refinancing. The "Adj B/E" column on the rate table shows the break-even including the tax adjustment and MI cost:

```
Adj B/E = (total_closing_costs + tax_adj + MI_diff) / monthly_payment_savings
```

If MI_diff is positive (refi MI costs more than existing MI), the break-even is longer. If MI_diff is negative (refi MI costs less, e.g., the borrower's LTV has improved), the break-even is shorter.

MI drop-off with appreciation: The home value appreciates at rate g per year, compounded monthly. As the home value rises and the balance falls, LTV improves. The month MI drops off (T) depends on both the amortization schedule and the appreciation rate. Faster appreciation = MI drops sooner. The default appreciation rate is 3% per year. This is not adjustable per client. In practice, using a moderately different appreciation rate (e.g., 2% or 4%) does not meaningfully change the MI calculation or the break-even. The difference in when MI drops off is typically only a few months.

Practical Notes

Financing closing costs vs. paying out of pocket:

The rate table treats closing costs the same way regardless of whether the borrower finances them or pays out of pocket. The system does not differentiate between the two for **refinance clients**. The reason is straightforward: the break-even suggestion is about whether the rate reduction justifies the cost of refinancing. That question has the same answer either way. The total cost is the same; only the payment method differs. Whether the borrower writes a check at closing or rolls it into the balance, the economic cost of refinancing is approximately the same.

For purchase clients, points are always treated as paid out of pocket. This is the correct assumption for purchase transactions — the borrower pays points at closing as part of the purchase transaction; they are not financed into the loan. The loan balance stays at the loan amount M. This means the monthly payment comparison in the purchase rate table is a pure rate-to-rate

comparison on the same balance. The break-even numbers reflect exactly how long it takes for the lower-rate payment savings to recover the upfront points cost.

What does change is the monthly payment. If the borrower finances $8,000 in closing costs, their new loan balance is higher, so their monthly payment is higher and their actual savings are roughly 10-15% less than what the table shows. Example on a $400,000 loan at 7% going to 5.75% with 2% closing costs ($8,000): paying cash saves $327/mo with a 25-month break-even. Financing the costs saves $280/mo with a 29-month break-even, a 4-month difference. The break-even suggestion doesn't meaningfully change either way. The monthly savings number will be lower if they finance, but the decision to refinance remains the same.

Why the suggestion is indifferent to financing: The ADL/Berger model that produces the refinance trigger doesn't take "amount bringing to table" as an input. It takes only the loan's fundamentals: balance, rate, term, fixed closing costs, tax rate, expected ownership, attention rate, and rate volatility. The output is a rate-drop threshold in basis points, and that threshold doesn't change based on how the borrower funds the closing costs.

When the system converts the bps threshold to a break-even in months, it has to use a real new payment, so the displayed months will shift slightly when ATB changes. But every row on the rate table uses the same formula structure (totalCosts divided by paymentSavings), so the suggestion target and the per-row BE shift together. The only reason they don't move in perfect lockstep is the points cost. A row with positive points has a higher total cost than the at-par target, so it absorbs more of the cash from ATB before maxing out. This means the gap between target and row can shift by a few months at most. That is justified: a row with points genuinely involves a bigger upfront cost, and the borrower has to recoup more before breaking even.

The underlying decision (the rate-drop threshold) is unchanged either way, so which rate meets the target stays the same. The system shouldn't take a position on whether the borrower finances or pays cash. That is a separate conversation between the LO and the borrower based on cash flow.

Also worth noting: financing $8,000 in closing costs on a 30-year loan costs roughly $16,800 in extra interest over the life of the loan, but whether to pay cash or finance is not a question BPS Hawk answers because it depends on the borrower's opportunity cost of capital. If the borrower's savings or investment rate is higher than the mortgage rate (e.g., earning 8% while the mortgage is 5.75%), financing the costs and keeping their cash invested is the smarter move. If their money is sitting in a low-yield account earning less than the mortgage rate, paying cash saves them the extra interest. That's a personal financial decision between the LO and the client. It varies by borrower and has nothing to do with whether the refinance itself makes sense. The system answers the rate question. How to pay for it is a separate conversation.

Recommended loan term for refinance:

I would suggest that when refinancing, the borrower should match the remaining term on their current loan. If they have 28 years left, they should refinance into a 28-year (or the closest available, e.g., 30-year) term. This keeps the comparison apples-to-apples. You're evaluating the rate improvement without introducing a term change that could mask or inflate the savings. Changing

the term is a separate decision from changing the rate, and mixing the two makes it harder to evaluate whether the refinance is worthwhile on its own merits.

Changing the new loan term does not change the rate table decision:

That said, the rate table compares rates at whatever new term the client selects. If the client changes the new loan term (e.g., from 30 years to 20 years), the payment amounts change, but the rate-to-rate comparison and the break-even suggestion still hold. The suggestion is about which rate reduction justifies the cost of refinancing. That threshold is the same regardless of the term chosen. The term decision is separate from the rate decision.

BPS Hawk does not tell the LO which term to choose:

The system is completely agnostic on loan term. It does not recommend a 30-year over a 15-year or vice versa. That decision depends on the client's cash flow, goals, and risk tolerance, things only the LO and borrower can evaluate. BPS Hawk's job is to say "at this rate, refinancing makes financial sense," not to say which loan product to take.

LLPA bracket changes from LTV or credit score improvements are not accounted for:

The system calculates LLPA based on the client's current LTV and credit score on file. It does not anticipate that these might change. For example, a client's home may only need to appreciate a small amount to cross into a better LTV bracket (e.g., from 80.5% to 79.9%), which would meaningfully reduce their LLPA. Similarly, a client's credit score might be just below a tier cutoff, and a small improvement could drop their LLPA significantly. BPS Hawk does not flag these edge cases or project future LTV/credit score improvements. The LO should be aware of clients who are close to bracket thresholds and update the property value or credit score on the client record when conditions change. This is an opportunity the system doesn't catch automatically.

The suggestion is a trigger, not a guarantee:

When the break-even meets the target, BPS Hawk is telling you it's time to look. The LO still needs to verify the numbers, check that the client's situation hasn't changed (life events, plans to move, credit changes), and have the conversation. The system does the monitoring and the math. The LO makes the final call with the client.