

Choose Texas Power: Product Recommendation

Kevin Middleton

Problem Framing

Tension between explainability and drop-off exists, but it is a symptom, not the root cause. The real problem is that customers do not trust the match.

Customers answered questions. They got recommendations. But they do not understand why these plans are right for them. Without that confidence, customers either drop off (too much friction to continue without clarity) or convert into regret, confused at purchase and frustrated when the first bill arrives. Both outcomes hurt the marketplace.

Deregulated energy is confusing by nature. Advertised rates do not match actual costs. Usage tiers, bill credits, base charges, and plan mechanics make comparison nearly impossible for most consumers. This is exactly why the intelligent matching experience exists. But if the AI does the work and the customer cannot see the reasoning, they are being asked to trust a black box in a category where trust is already low.

Why this problem matters more than others right now

The experience is only slightly contribution-margin positive, which tells me we are converting volume without earning confidence. It suggests the current flow is capturing orders but not building the kind of customer confidence that drives retention and longevity. REP incentives and partner mix shifts are real tensions, but they are downstream of customer trust. If customers trust the match, they stay longer. If they stay longer, REPs benefit from longevity incentives. If REPs benefit, the partner tension softens. **Trust is the leverage point.**

Recommendation

Make the AI's reasoning legible by showing customers their estimated actual monthly cost and the specific usage assumptions that make this plan the right fit.

This is not about adding more explanation. It is about making the explanation *worth* the friction by delivering an "aha moment" that builds confidence.

I have seen this work before. At [HVAC.com](https://www.hvac.com), I inherited a site with 2M annual visitors and a 0.1% conversion rate. The gap between "I'm researching" and "I'm ready for a quote" was too wide. I closed it by building [interactive calculators](#) that gave homeowners personalized answers: [repair or replace](#), [system sizing](#), [energy savings](#). Each tool delivered a moment of clarity that made the next step feel earned. Calculators peaked at ~7.3% conversion. [That is over a 50x lift from baseline](#).

[QuoteScore](#)¹ worked the same way: QuoteScore took something opaque (a contractor's quote) and broke down exactly what was good, what was bad, and what was unclear. When a score came back yellow or red, we offered to connect them with a fair quote from our network. Customers engaged because we helped them understand first, then offered a path forward.

The lesson was not "add more content." It was "give people something useful before asking for their commitment."

Value first, commitment second.

*Later, insights from that work informed [Trane's System Recommender Tool](#).² The value was directional guidance: hundreds of confusing choices narrowed to three that made sense for their home. The dealer could find a lookalike. You are not closing the sale online; you are **earning the next step**.*

Behavior change: Customers move from "I got a recommendation" to "I understand why this plan fits me and what tradeoffs I'm accepting." **That confidence improves conversion quality, not just conversion volume. Customers who understand their plan are less likely to churn after their first bill arrives.**

Trade-offs and Constraints

Optimizing for: Customer trust and retention, which I believe is the path to sustainable margin improvement.

Accepting: This may not maximize short-term order volume. Some customers who would have converted quickly may pause to absorb the additional context. I am betting that the customers who convert with confidence are worth more over time than the ones who convert confused.

This recommendation also does not address how REPs differentiate beyond price. Some providers offer better service, renewable options, or local support that may justify higher rates. A future iteration could give REPs a way to surface their value proposition alongside cost comparisons, so customers can make informed tradeoffs rather than defaulting to the cheapest rate.

¹ QuoteScore was sunset after HVAC.com was acquired by Trane, as it was deemed anticompetitive due to scoring HVAC quotes from competitor brands.

² System Recommender is currently live in NC, SC, FL, GA, TX, and AZ, with nationwide rollout planned.

Constraints that shaped this decision:

- **The core buy flow will not be reversed.** This recommendation works within the existing intelligent matching experience, not against it.
- **Leadership expects measurable progress in 1-2 quarters.** Showing estimated actual cost is achievable in that timeframe. It does not require rebuilding the recommendation engine, just surfacing its reasoning.
- **Internal stakeholders disagree on whether to optimize for growth, trust, or partner stability.** This recommendation picks trust, but frames it as the path to the other two. If I am wrong, the metrics will show it.
- **Regulatory scrutiny on estimated rates is increasing.** Proactively showing customers how we calculated their estimate is a defensive move as much as a product improvement.

How I Would Evaluate Success

Primary metrics:

1. **Conversion rate** (completion of order flow). Expecting a modest lift or hold, not a spike. The goal is quality, not volume.
2. **90-day customer retention.** If the hypothesis is correct, customers who understood their match should stay longer. This is the metric that matters most.
3. **REP longevity bonus payouts.** A lagging indicator, but the one that aligns partner incentives with customer outcomes.

What would cause me to revisit or reverse:

- If conversion drops materially (more than 10%) without a corresponding lift in retention within 90 days, the added context is creating friction without delivering value. At that point, I would revisit the *presentation* of the explanation, not scrap transparency. *I learned this at [Rocket Lawyer](#): we [hypothesized that hiding previews on shorter documents would improve conversion](#), but it backfired by 10%. The preview was building trust we had not accounted for. The lesson wasn't to stop iterating, but to dig into **what builds confidence**.*
- If this change does not improve retention, I am wrong about trust being the leverage point. But I would rather learn that from data than assume it away. **If we cannot explain why a plan fits, we have not earned the right to sell it.**

Approach Explanation

What I focused on:

- I framed this as a trust problem rather than an explanation problem. The three tensions in the brief (REP incentives, partner mix, explainability) are all real, but I think they share a common root: **customers converting without confidence creates downstream problems for everyone.** I focused on the lever I believe has the most compounding effect and the one I would have the most direct control over.
- I drew on my experience at HVAC.com and Trane, where I worked on similar marketplace dynamics: confusing product categories, consumers who need directional guidance rather than perfect precision, and what we called the Inequitable Power Dynamic, where the person you are seeking help from knows far more than you do and could take advantage of that gap. **Building trust is not optional in these environments.**

Tradeoffs I made:

- I picked one recommendation instead of hedging across multiple. The brief asked for that, but it also meant I could not address the REP incentive structure directly. If I had more space, I would have explored whether shifting compensation further toward longevity bonuses (rather than per-order payments) could reinforce the same trust-first approach from the partner side.
- I also kept the recommendation scoped to something achievable in 1-2 quarters. There are bigger bets available (rethinking the data collection flow, building post-purchase tools to reinforce the match), but those felt outside the constraint of "measurable progress" in the near term.

How I would improve with more time:

- **I would want to see funnel data:** where are customers dropping off, and what do exit surveys or support tickets say about why? My hypothesis is that the trust gap is the issue, but I would want to validate that before committing.
- **I would also want to talk to REPs.** The brief frames them as potentially adversarial to customer-centric matching, but I suspect some REPs would welcome better-fit customers if it meant lower churn and more longevity payouts. Understanding which partners are aligned and which are resistant would shape how to sequence the rollout.
- **I would stress-test the "estimated actual cost" recommendation with a few customers before building it.** The goal is an 'aha moment,' which I define as making the customer 10% smarter or giving them a realization that builds confidence to take the next step. But if the estimate feels too precise and then differs from reality, it could backfire. Getting the confidence interval right matters.