



The Credits Question

How Figma Make Users Weigh AI Allocation, Efficiency, and Control, and What They Will Actually Pay to Upgrade

A Choice-Based Study of 225 Current Figma Make Users.



Analyst & whitepaper editor: [Wesley Jacobsson](#) is a consumer insights expert who helps brands understand what drives purchasing decisions. His work has spanned sectors as varied as consumer goods, technology, travel, and finance. He specializes in advanced choice-based methods like conjoint and MaxDiff, translating complex behavioral data into insight brands can act on.

July 2026



Predicting Value.
Powering Decisions.

Executive summary

- **Make users are not rejecting paid AI; they will accept a higher price when the credit experience improves.** Raising the plan to \$30 with no other changes increases rejection. But pairing that price with spending control and efficiency features lifts acceptance from 64% to 79%. What earns the upgrade is predictability and control, not a larger pile of credits.
- **Credit allocation is the single largest driver of choice (28.4%), but control and efficiency together nearly match it (28.9%).** Users are not just asking for a bigger bucket; they place independent value on seeing, managing, and stretching the credits they already have.
- **The strongest tested configuration is the feature-and-price-optimized package (\$30):** the current 3,000-credit allocation plus usage alerts, per-prompt estimates, draft mode, and smart routing. The features offset the price resistance, cutting "none of these" among Full Seat users from 33% to 18%.
- **Control and efficiency outperform large credit tiers once AI cost is included.** On gross revenue, the higher-credit tiered portfolios look best, but every consumed credit carries inference cost. Under the base-case model, the feature-and-price-optimized package leads on AI-cost-adjusted revenue at \$1.54M, while the larger multi-package portfolios fall toward or below the current package.
- **Multi-direction generation is a feature, not the anchor.** It topped the earlier concept-ranking study on desirability but fell to last (5.8%) once it had to compete against price and credits.
- **The broader lesson for AI pricing:** monetization is won by making the meter understandable and predictable, not by hiding it or simply enlarging the bundle.

The credit frustration and the upgrade question

When Figma began strictly enforcing AI credit limits in March 2026, the [Figma community forums filled with complaints](#). Three themes recur. The allocation feels tight: Users have reported that the 3,000 monthly credits on the Professional plan equate to [roughly 50 to 70 Make prompts](#), which power users exhaust in a single working session. Consumption feels unpredictable, because an agentic prompt can quietly trigger many internal actions, so two similar requests can cost wildly different amounts. And it feels opaque: users report [no clear warning before the balance hits zero](#) and no per-action cost shown up front.

Crucially, the complaint is not that users refuse to pay. As one forum contributor put it, nobody is asking for free AI; they are asking for fair limits, fair pricing, and visible value. The useful question is not whether Make users are frustrated, which is well established, but what a plan built around their needs should actually contain: the combination of credit allocation, rollover, efficiency, control, and price they will pay for.

This matters because current Make users already sit inside the ecosystem, the overwhelming majority on the Professional Full Seat plan. The opportunity is not to win them as new customers but to identify the package configuration that justifies an upgrade to a higher tier. EPIC ran this study to answer that directly: to move past the well-documented frustration and quantify which features command real willingness to pay, where that willingness concentrates, and which complaints translate into upgrade revenue rather than just sentiment.

Building on what users told us they value

Before pricing those features, we needed to know which ones mattered most. An earlier EPIC concept-ranking study of 300 AI design-to-code users tested 22 features to establish what the category prioritizes when trade-offs are unavoidable. Its central finding was that users value tangible output over interaction mechanics, and one feature sat at the very top. Generating multiple design directions at once tied for the highest utility score, and the accompanying market audit found it to be one of the clearest white spaces in the category, supported as a native workflow by only a minority of platforms: a feature users value most that most platforms have not yet built well.

That made multi-direction generation the natural feature to carry into this pricing test. A high score in the concept-ranking study shows a feature is prioritized when users choose between capabilities at no cost to themselves; it does not show whether they will pay for it, or how much. This choice-based study puts that question directly: does the top-ranked feature command genuine willingness to pay among Figma Make users, or does its appeal soften once a price is attached and it must compete against simply having more credits? It is the



Predicting Value.
Powering Decisions.

only platform-capability feature carried into the design, as it was the only high-ranking feature that Figma Make lacks completely. Every other attribute is drawn from the credit economy users actually complain about.

Why this matters now: The credit question is not a narrow packaging issue. It sits at the center of how AI-native SaaS products will need to monetize usage. Figma's AI credits apply across a growing set of AI actions, including Make, image editing, and AI-powered workflows across the platform. As AI becomes more embedded in daily product use, the commercial challenge is no longer simply whether users value AI. They do. The harder question is how to price variable-cost AI usage in a way that feels fair, predictable, and controllable to the user while still protecting product economics.

What we tested

The study was built around a choice-based activity designed to estimate how Figma Make users trade off credit allocation, control, efficiency, capability, and price when those elements are bundled together in realistic plan options. Asking users in isolation whether they want more credits would tell us very little. The answer is obvious. The more useful question is what they are willing to give up, and what they are willing to pay, when more credits compete against rollover, usage transparency, efficiency tools, multi-direction generation, and monthly price.

The design centered on the AI-credit features users most directly experience and most often raise in community discussion, plus the one high-value capability carried over from the earlier concept-ranking study, multi-direction generation. Each plan concept was built from the following attributes and levels:

Attribute	Levels tested
AI credits included per month	3,000 / 6,000 / 9,000 / 12,000 / 25,000
Credit rollover	Expire at end of month / Roll over 1 month / Roll over 3 months / Roll over 6 months / Roll over 1 year
Price per credit beyond monthly allocation (pay-as-you-go)	Per-credit overage rates ranging from \$0.01 up to \$0.11
Spending control features	None / Automatic alerts at 50%, 75%, 90% of allocation / Per-prompt usage estimate based on similar prompts / Both alerts and per-prompt estimate
Credit efficiency features	None / Draft mode (lower-fidelity drafts at lower credit cost) / Smart routing (auto-selects the most efficient model per task) / Both draft mode and smart routing
Multi-direction generation	None (one output per prompt) / Parallelx2 (up to two directions per prompt, ~2x credit cost) / Parallelx4 (up to four directions per prompt, ~4x credit cost)
Total price per month	Monthly per-seat price points, scaled in steps so that higher credit allocations were tested against correspondingly higher prices (from roughly \$15 at the lowest tier up to the \$80s at the highest)

This design was built to answer a commercial question, not simply a preference question. The objective was not to confirm that users prefer more credits and lower prices. It was to understand which features carry enough weight to justify a higher-priced Make package, where willingness to pay concentrates, and which complaints translate into monetization potential rather than sentiment alone.



Predicting Value.
Powering Decisions.

Methodology

EPIC surveyed 225 current Figma Make users who use the platform at least multiple times per month. The sample was intentionally focused on users with direct experience of the product and its credit system. Approximately 81% of respondents were current Professional Full Seat users, with the remainder drawn from Free Starter, Collab seat, and Dev seat users.

This matters for interpretation. The study was not designed to represent the full universe of design-to-code users. It was designed to understand monetization potential among current Figma Make users, most of whom already sit inside the paid Figma ecosystem. The lower-seat users were included deliberately as an exploratory segment, to test whether any package configuration could also reduce friction to trading up into a Professional Full Seat-style offer.

All respondents were responsible for paying for their own Figma plans or making their own subscription decisions. This ensured that choices reflected personal willingness to pay rather than passive usage of an employer-paid account.

By asking respondents to choose the plan they would actually pay for, or reject the options entirely, the choice-based study allowed us to observe how package specification and price interact. This moves beyond simple stated preference. Instead of asking whether users value transparency, efficiency, or more credits, it shows how much those elements matter when users must choose between them.

What the study reveals

When we modeled the trade-offs across all seven attributes, one result stood out clearly: the number of AI credits included in the plan was the single largest driver of choice, with 28.4% relative importance.

Price followed at 17.1%, then spending control at 15.5% and credit efficiency at 13.4%. Pay-as-you-go overage pricing came next at 11.1%, ahead of credit rollover at 8.7%. Multi-direction generation, the feature carried over from the earlier study, landed last at 5.8%.




Figure 1. Relative feature importance.

The headline finding is intuitive, but the structure underneath it is more commercially useful. Users do care deeply about the size of their credit allocation. That is the anchor of the decision. But demand is not only about more credits and lower prices. Spending control and credit efficiency together account for 28.9% of relative importance, effectively matching the importance of raw credit allocation itself.

That answers one of the central questions of the study. Users are not simply asking for a larger bucket. They place real, independent value on being able to see, control, and stretch the credits they already have.

The surprise sits at the other end of the ranking. Multi-direction generation tied for the highest utility score in the earlier study, where users evaluated feature desirability without price attached. In this choice-based study, where it had to compete against credits, controls, efficiency, and monthly cost, it fell to the bottom of the ranking.

That does not mean multi-direction generation has no product value. It does mean it should not be the lead paid differentiator for an upgraded package. Once price enters the decision, users prioritize the economics of the credit system itself: how many credits they get, what they cost, whether usage is predictable, and whether they can manage consumption without surprises.

A higher-priced package built primarily around a marquee capability would therefore be anchored in the weakest attribute in the study. A package built around a larger, clearer, better-managed credit experience is aligned with the strongest ones.

What users prefer within each attribute

Relative importance shows how much each attribute moves choice overall. Looking inside each attribute at the preference score for each level shows which specific configurations create or suppress demand.




Figure 2. Relative feature-level preference.

Three patterns are commercially important.

First, the value of credits is real, but not unlimited within the tested price ranges. Preference rises from 13.0 at 3,000 credits to a peak of 14.9 at 6,000 credits and 14.8 at 9,000 credits. It then eases back, with the 25,000-credit tier scoring lowest at 11.1.

This should not be read as users actively disliking credits. Because larger allocations were tested at higher prices, the dip most likely reflects the price burden attached to very large bundles. The practical signal is that preference concentrates in the mid-range, where users see a meaningful improvement in allocation before the price premium becomes too steep.

Second, within spending control, preference rises from 11.1 when no controls are included to 15.9 when automatic alerts and per-prompt estimates are offered together. That is a clear signal: users do not just want more credits. They want to know where those credits are going before they disappear.

Efficiency features follow the same direction, although in a tighter band. Draft mode, smart routing, and the combined feature bundle all perform well above the no-feature baseline, but users distinguish less sharply between the specific mechanisms. The commercial takeaway is that the presence of credible efficiency tooling matters more than the exact form it takes.

Third, the lower-importance attributes show flatter level curves. Any rollover is preferred to end-of-month expiry, but rollover options cluster closely together. Pay-as-you-go rates vary only modestly in preference across the tested range, suggesting that users focus more on the monthly allowance and plan experience than on the precise overage rate. Multi-direction generation improves preference only modestly, from 11.1 with no multi-direction generation to 13.0 for Parallel x4.

Price behaves as expected. Preference declines smoothly as the monthly price increases, from 16.0 at the base price level to 11.1 at base plus 100%. There is no sharp cliff in the tested range. That suggests this user base will tolerate higher prices when the package feels sufficiently better.

Commercial read: Rollover is appealing because it addresses perceived waste. But the choice-based research shows that rollover is not the strongest lever. Spending control and efficiency matter more because they address the problem users experience in the moment. Rollover compensates users after the fact. Control and efficiency help prevent the problem before credits are spent.



Predicting Value.
Powering Decisions.

From preference to package economics

The choice-based study results show what users value. The next question is what those preferences are worth commercially.

The simulations below convert choice-based utilities into predicted selection shares for specific plan configurations, applied to an illustrative base of 100,000 current Figma Make users. The 100,000-user base is an indexing device. It is not an estimate of Figma's actual Make user population.

Credit rollover is held at one month and pay-as-you-go pricing at \$0.03 per credit throughout the main simulations. Both carry lower relative importance than credit allocation, price, spending control, and efficiency, and their level curves are relatively flat within the tested ranges, so holding them fixed isolates the levers that move choice most.

How to read “none of these”

The “none of these” option, the study's outside option, is essential to the analysis, but it must be interpreted carefully.

For the purposes of comparing tested packages, respondents selecting “none” are assigned no revenue from the tested configuration. That is appropriate for understanding which package captures the strongest active choice. However, it should not be read as churn, cancellation, or zero total customer revenue to Figma. Crucially, what “none” becomes in the real world depends on the choice architecture: an optional plan switch likely means users stay put, whereas a forced migration could turn some of that rejection into cancellation.

This is especially important because the sample is concentrated among current Professional Full Seat users. For these users, “none” most likely means non-acceptance of the tested upgrade or repriced package. They may remain on their current plan, continue paying for their existing seat, or reject the specific package shown. For Free Starter, Collab seat, and Dev seat users, “none” is closer to non-upgrade into a Professional Full Seat-style offer.

The single-plan revenue view

The natural starting point is the current package: 3,000 credits at \$16 per month (billed annually). In the simulation, this baseline captures roughly 71% of users and produces \$1.14 million in selected package revenue against the illustrative 100,000-user base.

The first lever is price. The price elasticity curve for the current package shows room to move. Raising the single plan from \$16 to \$30 reduces selection, but not enough to offset the higher price. Predicted selection falls to roughly 64%, while selected package revenue rises to \$1.92 million.

Interpretation note: Throughout this paper, selected package revenue means simulated revenue from the packages respondents actively choose within the choice-based model. It is a disciplined way to compare package configurations, prices, and trade-offs. It is not the same as an observed booking or subscription forecast. Actual outcomes will depend on live product usage, billing architecture, admin purchasing behavior, competitive context, seasonality, and Figma's own cost structure. The results should be read as directional commercial signals to validate in market, not as guaranteed demand.

Note: \$30 was the highest price tested for the 3,000-credit configuration. The result therefore shows that the 3,000-credit package can support a materially higher tested price than today's \$16 plan, particularly when paired with control and efficiency features. It does not establish \$30 as the absolute revenue-maximizing price or the upper limit of willingness to pay.




Figure 3. Price elasticity curve for the baseline plan.

The second lever is the control and efficiency feature set. Holding price at the highest tested 3,000-credit price point of \$30, and adding the strongest tested levels of spending control and efficiency, automatic alerts plus per-prompt usage estimates, and draft mode plus smart routing, lifts selection from 64% to 79%. Selected package revenue rises to \$2.37 million.




Figure 4. Simulated monthly package revenue after price and feature optimization.

This is a critical result. The features materially offset the resistance created by the higher price. The price-optimized package (\$30) is easier for users to reject. The feature-and-price-optimized package, at the same \$30, feels more justifiable, because it directly addresses the pain points users associate with the credit system.



Predicting Value.
Powering Decisions.

Because the sample consists of current Figma Make users, this should not be read as new-user acquisition. It is better understood as stronger monetization acceptance within the existing Make user base: fewer respondents choose “none of these” when the higher-priced package includes clearer controls and better credit efficiency.

Who the optimized package moves

Before moving from single-plan optimization into tiered portfolios, we cut the same three single-plan simulations by current seat type. The purpose was to test whether the optimized \$30 package is only a monetization lever for current Professional Full Seat users, or whether it also makes a Professional Full Seat-style offer more compelling to users currently on Free Starter, Collab, or Dev seats.

Segment	Current package (\$16): “none of these”	Price-optimized package (\$30): “none of these”	Feature + price-optimized package (\$30): “none of these”	“None of these” reduction after feature optimization
Full Seat users	26.45%	33.49%	18.16%	15.33 pts
Free / Collab / Dev users	40.26%	46.86%	33.18%	13.68 pts

Table 1. “None of these” share by current seat type across single-plan simulations.




Figure 5. “None of these” share by current seat type.

The pattern is consistent across both groups. Raising the price of the current package from \$16 to the price-optimized package at \$30 increases “none of these” selection. Among Full Seat users, “none of these” rises from 26.45% to 33.49%. Among Free, Collab, and Dev users, it rises from 40.26% to 46.86%. That is the price penalty. A higher monthly price captures more revenue among those who select it, but the price-optimized package is easier to reject.

Adding spending-control and credit-efficiency features more than offsets that resistance. Among Full Seat users, “none of these” share falls from 33.49% for the price-optimized package to 18.16% for the feature-and-price-optimized package. Among lower-seat users, it falls from 46.86% to 33.18%.



Predicting Value.
Powering Decisions.

In both groups, the feature-and-price-optimized package performs better than the current package on “none of these” selection. That matters. The features do not simply defend a higher price. They make the higher-priced package more attractive than the current package in an active choice setting.

The implication is not that the optimized package is mainly a lower-seat upgrade lever. Full Seat users remain more receptive in absolute terms, with 81.84% selecting the feature-and-price-optimized package versus 66.82% among lower-seat users. The stronger reading is that the same package plays two roles. It supports price acceptance among current Full Seat users and provides directional evidence of trade-up appeal among users not currently on a Professional Full Seat.

This strengthens the recommendation. Control and efficiency features are not cosmetic product improvements. They make a higher-priced credit package easier to justify. The effect is clearest among users already on the Full Seat plan, which keeps the core story anchored in monetization of the existing paid base, but the lower-seat response points to the fact that even those on lower tiers ascribe value to these features.

Why test portfolios

The subgroup cut strengthens the case for the optimized single plan, but it does not remove the structural limitation of a one-package strategy. Even after feature optimization, some users still choose “none of these”. Those users are not all rejecting the offer for the same reason. Some may need a lower-friction entry point. Others may need a larger allocation or a clearer upgrade path.

That raises the portfolio question: can a tiered lineup capture demand that a single plan leaves behind?

We tested several potential three-package lineups. Each ladder two things together across its tiers: credit allocation and multi-direction generation.

Tier	Credit allocation	Multi-direction generation
Base	3,000 credits	None
Middle	6,000 or 9,000 credits	Parallel x2
Highest	9,000, 12,000, or 25,000 credits, depending on lineup	Parallel x4

On selected package revenue, the portfolios perform strongly. Offering three plans rather than one cuts the “none of these” rate sharply. The strongest gross-revenue lineup, 3,000 / 9,000 / 25,000 credits, reaches \$3.36 million, nearly three times the current package’s \$1.14 million.

On package revenue capture alone, the portfolios win cleanly. They give users more ways to find a plan that fits, and the highest-allocation lineups generate the largest gross revenue. But gross revenue is not the whole story in an AI-credit business.

The AI-cost-adjusted view

Every consumed credit carries an inference cost. To estimate the economics of included credit allowances, we applied a base-case AI cost model to the simulated package revenue.

For this model, we estimate cost at \$0.014 per consumed credit. This is derived from Figma’s published \$0.03 pay-as-you-go overage rate and an [external benchmark for AI-product gross margins of approximately 50%](#). This is not Figma-disclosed cost data. It should therefore be read as a base-case modeling assumption, not as a statement of Figma’s actual unit economics.

Because most subscribers are unlikely to consume their full allowance every month, we apply a 25% utilization rate to included credits. The effective cost is therefore \$0.0035 per allocated credit. AI-cost-adjusted package revenue is calculated as selected subscription revenue minus estimated AI credit cost. This blended 25% is applied uniformly across tiers for simplicity. In practice, utilization is unlikely to be flat: the credit-constrained users who exhaust the current 3,000-credit allowance in a single session may run well above 25%, which would



Predicting Value.
Powering Decisions.

raise the cost attached to the smaller tiers, while large-tier buyers tend to self-select for heavier use. The tier ranking that follows is therefore conditional on this assumption.

This is not full net revenue and not a complete P&L. It excludes non-AI product costs, support, sales, marketing, R&D, payment fees, and broader platform economics. It is a focused view of how selected package revenue changes once the cost of included AI credits is taken into account.

Under those assumptions, the ranking changes materially.

Configuration	Gross revenue	AI credit cost	AI-cost-adjusted package revenue	vs. baseline
Feature + price-optimized (\$30)	\$2,373,108	\$830,550	\$1,542,558	+\$1,151,719
Price-optimized (\$30)	\$1,918,401	\$671,475	\$1,246,926	+\$856,087
Portfolio 3 / 6 / 9	\$2,682,061	\$2,013,480	\$668,581	+\$277,742
Portfolio 3 / 6 / 12	\$2,832,255	\$2,241,540	\$590,715	+\$199,876
Portfolio 3 / 9 / 12	\$3,028,543	\$2,533,230	\$495,313	+\$104,474
Current package (\$16)	\$1,137,074	\$746,235	\$390,839	—
Portfolio 3 / 9 / 25	\$3,357,989	\$3,039,575	\$318,414	-\$72,425
Portfolio 3 / 6 / 25	\$2,884,549	\$3,405,255	-\$520,706	-\$911,545

Table 2. Selected package revenue, estimated AI credit cost, and AI-cost-adjusted package revenue by configuration. Illustrative 100,000-user base. Assumes \$0.014 per consumed credit and 25% utilization of included credits.



Predicting Value.
Powering Decisions.

The two optimized single-plan configurations, which sit behind the portfolios on selected package revenue, finish first and second once estimated AI credit cost is included. The feature-and-price-optimized package is the clear leader, with \$1.54 million in AI-cost-adjusted package revenue, compared with \$390,839 for the current package.

The 3 / 6 / 9 portfolio improves materially on the larger-credit portfolios and is the strongest tiered option tested on an AI-cost-adjusted basis. Even so, it remains well below the feature-and-price-optimized single-option lineup. The issue is not that portfolios fail to attract users. They do. The issue is that the larger included credit allowances add cost faster than they add monetizable willingness to pay.

The revenue-leading 3 / 9 / 25 portfolio illustrates the trade-off most clearly. It generates the highest selected package revenue, but under the assumed cost model it falls below the current package on AI-cost-adjusted package revenue. The 3 / 6 / 25 lineup performs worse, turning negative after estimated AI credit cost.

Sensitivity note: The AI-cost-adjusted ranking should be read as a base-case view, not a fixed truth. The conclusion is most sensitive to two variables: cost per consumed credit and utilization of included credits. If higher-credit subscribers consume materially less than the assumed 25% of their allowance, larger tiers become more viable. If they are heavy users who consume more than the blended assumption, the economics deteriorate quickly. This is why the recommendation is not to abandon tiering altogether, but to validate larger allowances through usage telemetry before building the commercial strategy around them.

Variable to validate	Why it matters	Commercial implication
Cost per consumed credit	The \$0.014 figure is a derived base case, not disclosed by Figma, and it drives the entire AI-cost-adjusted ranking.	Confirm against real inference cost before committing to credit-heavy tiers.
Utilization of included credits	A flat 25% is assumed across all tiers. Heavy large-tier users would raise cost; lighter users would lower it.	Measure actual consumption by credit tier and seat type through usage telemetry or a controlled pilot.



Predicting Value.
Powering Decisions.

Why the larger tiers struggle

The source of the inversion is visible when contribution is split by included credit tier.




Figure 6. Average contribution to AI-cost-adjusted package revenue by included credit tier

Across the simulated portfolios, the 3,000-, 6,000-, and 9,000-credit tiers contribute positively under the base-case cost model. The 12,000- and 25,000-credit tiers do not. The issue is not that users place no value on larger allowances. They do. The issue is that willingness to pay does not rise quickly enough to offset the assumed cost of servicing those larger allowances at 25% utilization.

A separate willingness-to-pay exercise provides additional directional support. When respondents were asked directly about an unlimited-credit package that was otherwise equivalent to the current Figma Professional Full Seat, willingness to pay did not rise to a level that would clearly support very large included allowances. Acceptance weakened materially above \$44 per month. That reinforces the core pattern from the choice-based study: willingness to pay for incremental credits appears to flatten as packages get larger.

This should be treated as an economic signal, not a final pricing answer. The true viability of larger tiers depends on real usage behavior. If buyers of higher-credit packages consume far less than their allowance, the economics improve. If they are heavy users who consume materially more than the blended average, the economics deteriorate.

What this means for the lineup

The strongest commercial signal is not “sell more credits.” It is “make the credit economy easier to manage.”

Users value allocation, but they also value predictability, control, and efficiency. The most attractive package is not the one that simply offers the largest credit pool. It is the one that gives users confidence that they can understand, manage, and stretch the credits they have.

Under the base-case economics, the single most profitable configuration tested is the \$30 plan with the current 3,000-credit allocation, automatic usage alerts, per-prompt estimates, draft mode, and smart routing. This package does four things at once:

- It creates a clear reason to accept a higher monthly price.



Predicting Value.
Powering Decisions.

- It directly answers the pain points users raise around opacity, unpredictability, and wasted credits.
- It limits the cost exposure created by very large included allowances.
- It leaves multi-direction generation as optional upside, not the anchor.

The final recommendation

The final recommendation is clear: lead with control, efficiency, and confidence. Figma Make users are not rejecting paid AI. They are rejecting a credit experience that feels too tight, too hard to predict, and too difficult to manage.

Under the base-case assumptions, the strongest tested commercial configuration is a disciplined \$30 package that keeps the 3,000-credit allocation but adds the features users need to make that allocation feel visible, controllable, and efficient. Larger included-credit tiers may still have a role, but they should be validated against real utilization before becoming the center of the pricing architecture.

Finally, this study was designed to understand monetization potential within Figma's existing user base. A natural follow-on question is whether the same package has competitive reach beyond it. Testing the optimized configuration against users of competing platforms would show whether control-and-efficiency packaging can attract switchers, expanding their overall base, rather than just upgrade existing users.

The broader lesson for AI pricing

Figma Make is a specific case, but the lesson is broader. As AI features become embedded in SaaS products, companies face a new pricing challenge: users want the freedom of unlimited usage, while providers need to manage variable inference costs.

The answer is not simply stricter limits or larger bundles. The answer is a better value exchange.

AI monetization will not be won by hiding the meter. It will be won by making the meter understandable, predictable, and worth paying for.

Key takeaways

- 1. The credit economy is the product experience.** Users are not just evaluating AI output. They are evaluating whether the credit system feels understandable, fair, and controllable.
- 2. Allocation leads, but control and efficiency are almost as important.** Credit allocation is the largest single driver of choice, but spending control and credit efficiency together effectively match its importance.
- 3. Multi-direction generation is a feature, not the paid-tier anchor.** It performed strongly in the concept-ranking study, but fell to the bottom of the choice-based study once price and credit trade-offs were introduced.
- 4. The \$30 control-and-efficiency plan is the strongest tested single-plan direction.** It raises acceptance versus the price-optimized package and delivers the strongest AI-cost-adjusted package revenue under the stated assumptions.
- 5. "None" is not churn.** For current Professional Full Seat users, "none of these" means non-acceptance of the tested package, not zero customer revenue to Figma.
- 6. Larger tiers need usage validation.** Large included allowances can win selected package revenue but become economically sensitive once AI credit cost and utilization are included.
- 7. The broader AI pricing lesson is control.** Users are more willing to pay for AI when the meter is visible, predictable, and connected to clear value.

How EPIC can help SaaS leaders price AI with confidence

This is not just a Figma Make question. It is the monetization challenge facing every SaaS business building AI into the product experience. When usage creates variable cost, pricing decisions need to balance adoption, willingness to pay, margin protection, and user trust.



Predicting Value.
Powering Decisions.

EPIC helps product, pricing, growth, and commercial teams move beyond opinion by simulating the choices users would make in market. Using choice-based research, pricing simulation, segmentation, and cost-adjusted revenue modeling, EPIC can:

- **Quantify willingness to pay** for AI features, credit bundles, usage tiers, add-ons, and upgrade paths.
- **Identify which features justify an upgrade** and which should be metered, bundled, or treated as product experience improvements.
- **Model plan selection**, “none of these”, **trade-up**, **cannibalization**, and **portfolio revenue** across alternative package lineups.
- **Segment users by behavior, usage pressure, and price elasticity** rather than broad demographics alone.
- **Stress-test AI monetization against unit economics** including cost per task, credit utilization, overage pricing, and margin assumptions.
- **Translate the evidence into a board-ready recommendation** for pricing, packaging, product design, and market validation.

The result is not just a preference readout. It is a commercial decision framework that shows what users will choose, what they will pay for, where margin risk sits, and which package architecture is most likely to convert demand into profitable growth.

Talk to EPIC

If your team is deciding how to price AI features, credit bundles, usage tiers, or upgrade paths, EPIC can help you test the trade-offs before the market does. We turn complex pricing and packaging questions into clear commercial direction, showing what customers will choose, what they will pay for, and where the economics create risk or opportunity.



Predicting Value. Powering Decisions.

Pricing isn't static, it's a competitive weapon. If you are a brand that wants to quantify the equity and value your brand commands, alongside the willingness to pay for your products' features to create the optimum combination of product features and price, we'd love to hear from you.

EPIC Insights specializes in providing cost-effective, expert-led product and pricing conjoint software that delivers actionable retail, consumer products and services insights in just **14 business days**. By combining a cutting-edge research platform with deep domain expertise, EPIC empowers global leaders like **Coca-Cola, L'Oréal, Kraft Heinz, PepsiCo, ADP, Pleo, Square, DreamHost and Zenfolio** to make precise, data-driven commercial decisions at scale.

To meet with one of our experts visit

EPICINSIGHTS.IO