Balancing Customer Engagement and Annoyance in Online Retail: Insights from a Field Experiment

Jackie Baek
Leonard N. Stern School of Business, New York University, New York, baek@stern.nyu.edu

Daniel Chen
Carroll School of Management, Boston College, Boston, daniel.chen@bc.edu

Will Ma
Graduate School of Business, Columbia University, New York, wm2428@gsb.columbia.edu

Dmitry Mitrofanov
Carroll School of Management, Boston College, Boston, dmitry.mitrofanov@bc.edu

The emergence of e-commerce and digital operations has given firms unprecedented reach: retailers can now send real-time emails and mobile notifications to customers anywhere in the world. These direct-to-consumer promotions (e.g., coupons, emails, push alerts) can boost engagement and sales, but when overused, they can create annoyance and notification fatigue, alienate customers, increase unsubscriptions, and ultimately weaken long-term customer spending. In this paper, we first conduct a large-scale randomized field experiment at an online retailer to study this trade-off between engagement and annoyance. We find that sending fewer emails significantly reduces unsubscription (proxy for churn) by 59% and thus increases future customer survival, but at the cost of a 5-8% decrease in short-term revenue and purchase rate. These effects are heterogeneous: reducing coupon-email frequency lowers churn, especially among customers who previously used few coupons but opened many emails, and it reduces short-run spending, particularly among higher-volume customers and those who rarely opened emails. Then, using historical data over a longer time frame, we find that unsubscribing has a large negative effect on monthly revenue, with a 36% decrease in customer spending. We then embed these results in a customer lifetime-value (CLV) optimization model. The model shows that personalizing email frequency using the field-experimental estimates can improve long-term value: a myopic policy (maximizing immediate revenue) raises average CLV by 6.8%, while a full CLV-optimization policy yields an 8.3% gain over the daily-email baseline. These gains are driven by higher retention as well as modest changes in coupon usage. Overall, our findings highlight the importance of balancing short-term revenue with customer retention: sending emails with coupons too frequently can boost immediate sales but may "burn out" customers, whereas moderately reducing email frequency can yield higher lifetime value in many settings.

Key words: field experiment, customer annoyance, email frequency, customer lifetime value, online operations.

1. Introduction

As retail has become increasingly digitized and the emergence of e-commerce and digital operations has accelerated, firms have gained the ability to communicate with customers instantly and at high frequency through email, mobile apps, and social media. Many retailers now rely on automated newsletters, push notifications, and triggered messages (e.g., cart reminders, price-drop alerts, personalized coupons/emails) to bring shoppers back to the platform. In principle, these digital nudges can stimulate repeat purchases and bolster short-run sales because the marginal cost of sending an additional message is low and experimentation is easy. Yet the same scalability creates a fundamental tension: when promotions arrive too often or feel poorly targeted, customers experience annoyance and message fatigue, begin to ignore communications, unsubscribe from email lists, disable notifications, or delete the app altogether. A prominent illustration comes from Groupon. In an interview reflecting on the company's growth, cofounder Andrew Mason described internal pressure to increase deal and email frequency, recalling his initial reaction to the idea of sending two emails per day as "That sounds awful. Who wants to get two emails every single day from a company" (Blumberg 2018). He nevertheless agreed to run controlled tests, and the short-run results looked favorable: unsubscribes rose only slightly while incremental purchases more than offset the immediate attrition, making higher frequency appear "rational" by near-term metrics (Blumberg 2018). Mason's retrospective takeaway, however, was that hypergrowth left little time to observe longer-run churn dynamics, so the cumulative erosion of goodwill surfaced later, after repeated exposure had already trained customers to tune out or opt out; in his words, the service eventually became "a vestige of what it once was" (Blumberg 2018). The Groupon episode highlights why optimizing promotional intensity based solely on short-run lift can be misleading when annoyance and fatigue compound over time.

Similar tensions have surfaced well beyond Groupon. In the broader daily-deals ecosystem, it was noted that the very ease of "broadcasting" promotions at scale contributed to consumer annoyance and inbox saturation, with reported email response rates falling sharply as customers became desensitized to constant deal blasts (Schonfeld 2011). Related e-commerce models underscore how deeply frequent outreach can become embedded in operating strategy. For example, flash-sale retailer Zulily reportedly experienced meaningful revenue headwinds when deliverability problems reduced its ability to send daily marketing emails, illustrating both the power of high-frequency communication for short-run sales and the fragility of a model that relies on constant contact (Cook 2014). At the same time, firms have increasingly treated message volume itself as a design lever to manage annoyance and protect long-run engagement. Fab, the design e-commerce company, was reported to proactively reduce emails to users who were not opening them, and LinkedIn publicly described cutting a substantial share of its outgoing emails in response to member complaints, highlighting an emerging managerial consensus that the incremental lift from additional messages must be weighed against the cumulative cost of fatigue and opt-outs (Weber 2013, Awan 2015).

This modern context, when ubiquitous digital promotion is accompanied by the risk of overwhelming customers, motivates our study. We examine the engagement-annoyance tradeoff in online retail by leveraging both a randomized experiment and observational data from a retailer.[1] In a field experiment, we randomly assigned customers to receive either the status quo daily email frequency or a reduced every other day email frequency for a week. We measure the impact of this frequency decision on short-term outcomes (e.g., daily revenue, purchase rates) as well as on churn (unsubscribe) rates. Separately, using the retailer's historical data, we quantify the long-term consequences of unsubscribing on customers' future purchases and retention. Finally, we integrate these findings into a CLV optimization framework. Using the estimated effects of email frequency on short-term revenue and churn (from the field experiment) together with the long-term effects of churn on revenue (from the historical analysis), we simulate optimal policies for sending emails and coupons. We compare three policies: (1) the current baseline (daily emails to all), (2) a myopic policy that maximizes immediate expected revenue by tailoring email frequency, and (3) a full CLV-optimal policy that maximizes expected lifetime revenue accounting for retention.

Our main findings are threefold. First, the field experiment confirms a clear tradeoff: sending more emails per week significantly increases daily sales and purchase probability but also substantially raises the probability of customers unsubscribing (i.e., opting out of the email list). In other words, the more aggressive promotional policy drives short-term revenue but at the cost of losing some customers. These effects are heterogeneous. Reducing the frequency of coupon emails mitigates churn and unsubscribing, with the largest improvements among newer customers and those who historically used fewer coupons and tended to open emails more often. At the same time, sending fewer coupon emails lowers short-run spending, particularly among customers with higher past purchase volume and those who historically did not open many emails. Second, our analysis of historical panel data demonstrates that customers who unsubscribe from promotional emails suffer large drops in spending thereafter. Using regression and IV methods, we find that unsubscribing reduces customers' monthly revenue by a significant amount and markedly lowers customers' retention probability. Third, by plugging these estimates into a CLV optimization framework, we show that an optimized, personalized email policy can improve the firm's CLV. In an unconstrained setting (no budget limit on coupons), the myopic policy boosts average CLV by 6.8% over baseline, while the full optimization policy that maximizes expected lifetime revenue accounting for retention yields a 8.3% improvement. With a coupon usage budget fixed to status quo levels, the gains are smaller but still significantly positive (CLV +4.5%). Finally, we conduct a sensitivity analysis: as the base customer survival rate increases (i.e., when customers are generally more sticky), the

[1] Under our NDA, we cannot disclose the retailer's name; however, it is a major North American e-commerce retailer.

relative advantage of reducing email frequency grows significantly. Intuitively, when retention is already high, protecting that retention with less aggressive emailing yields larger lifetime revenue gains.

Overall, our contributions are both empirical and methodological. Empirically, we provide (to our knowledge) the first large-scale field experiment quantifying the short-term impact of email frequency in a retail setting as well as its long-term implications, including heterogeneous effects by customer type. Methodologically, we show how to combine field and observational data into an optimization of CLV, highlighting how direct causal estimates of engagement and churn can inform CRM decisions. Our findings also resonate with recent literature on advertising annoyance and marketing avoidance: for example, Todri et al. (2020) find that overly frequent online ads can reduce downstream purchase intent, illustrating a similar trade-off but without observing annoyance explicitly and modeling it through latent variables. In our setting, the ability to unsubscribe acts as a strong behavioral outcome of such annoyance, and this measure is fully observable to us. By explicitly measuring how frequency affects both immediate sales, annoyance, and CLV, we shed light on the optimal design of email campaigns to maximize long-term value.

Importantly, a unique aspect of our setting and dataset is that we observe a direct and explicit signal of customer annoyance through email unsubscriptions. This allows us to account for annoyance not only through subsequent purchasing behavior, which is a much more indirect proxy, but also through a clear behavioral action that reflects disengagement. This distinction differentiates our paper from much of the existing literature. For example, Todri et al. (2020) must infer annoyance implicitly using a hidden Markov model, which introduces additional identification challenges and potential bias.

The remainder of the paper is organized as follows. Section 2 reviews related literature on sales-annoyance tradeoffs, promotional/coupon field experiments, promotion/coupon optimization, and customer lifetime value and personalization. Section 3 describes our data: the retailer's historical customer panel and the randomized email frequency trial. Section 4 presents the results: first from the field experiment, then from the long-term behavioral analysis. Section 5 develops the CLV optimization model, reports counterfactuals of various email policies, and provides a discussion and managerial implications. Section 6 concludes.

2. Related Literature

Our study relates to four research streams spanning operations, marketing, and data-driven decision making. We build on work on engagement-annoyance tradeoffs, showing how higher contact frequency can lift short-run sales while increasing opt-outs and disengagement. We contribute new evidence from a large-scale coupon field experiment that tracks both immediate spending and downstream outcomes such as unsubscribing and repeat purchasing. We connect to promotion optimization by treating email frequency as the policy lever. Finally, we advance CLV and personalization by using heterogeneous treatment effects to design retention-aware, segment-specific email communication policies that balance short-run revenue with long-run customer value.

2.1. Engagement-Annoyance Tradeoffs

Frequent promotional communications can boost short-term sales but also risk annoying customers, leading to adverse long-term outcomes such as disengagement or churn. In online advertising, for example, high-frequency ad exposures have been shown to trigger consumer annoyance and avoidance. Todri et al. (2020) capture this tension, finding that while repeated display ads move consumers down the purchase funnel, excessive exposures can backfire by increasing irritation and reducing purchase intent. In direct marketing, firms face a similar tradeoff between revenue and customer goodwill. Early direct-mail models recognized that sending too many messages leads to diminishing returns in customer response. For instance, Bitran and Mondschein (1996) model diminishing returns to repeated mailings, assuming later mailings within a campaign yield lower response rates and profitability. More recent work in service operations shows that aggressive cross-selling promotions can impair service quality, effectively "annoying" customers in operational settings. Armony and Gurvich (2010) demonstrate that promotional offers during service calls lengthen call times and increase customer wait times, forcing managers to trade off incremental sales against deteriorating service metrics. Our study directly quantifies this sales-annoyance tradeoff in a field experiment by varying email frequency. By tracking both short-run spending and opt-out (unsubscribe) behavior, we provide empirical evidence of how a more aggressive email policy drives immediate revenue at the cost of higher customer annoyance (measured via unsubscribing). This extends prior work by explicitly measuring the customer annoyance reflected by unsubscription rate and the long-term revenue impact of annoyance, an aspect often modeled only indirectly in advertising studies (e.g., Todri et al. 2020). In doing so, we shed light on how firms can adjust contact frequency to manage the delicate balance between engagement and irritation in digital retail.

2.2. Promotional/Coupon Field Experiments

A rich stream of field experiments has examined the effectiveness of promotions and coupons on customer behavior. These experiments typically establish causal effects of promotional incentives on purchase incidence, basket size, or retention. For example, Sahni et al. (2017) run 70 field experiments with targeted discounts and found that beyond immediate uptake, targeted offers can act as advertising by attracting new customers and increasing overall store traffic. Similarly, large-scale field tests of coupons have revealed nuanced effects: Reimers and Xie (2019) show that offering deep e-coupons boosts short-term demand and can price-discriminate to attract new buyers, though the long-run retention of these customers is modest. Relatedly, Baek et al. (2025) study daily personalized coupon assignment for a retailer with over 20 million customers, proposing a scalable allocation policy that satisfies global discount business constraints and delivers a 4.5% revenue lift in an A/B test. Then, the authors further introduce a simple intertemporal reference-effects model in which customers benchmark against the best promotion seen over the past ℓ days, yielding an optimal cyclic structure that alternates several "worse" promotions with an occasional "better" one. In a similar spirit to our work, Wang et al. (2022) examine how email engagement, measured by open and click rates, relates to long-term profitability in a subscription retail setting, using a large-scale field experiment to uncover heterogeneous effects. While their study offers important insights into engagement–profitability links, it does not consider how email frequency influences opt-out behavior or churn. Our paper complements and extends their work by directly modeling the retention consequences of over-communication and optimizing frequency to balance engagement and annoyance in pursuit of higher customer lifetime value. Other experiments highlight how promotion design influences customer activity. For example, Kadiyala et al. (2024) run a retailer field experiment comparing delayed incentives (post-purchase gift cards) to instant discounts and find that delayed rewards boost engagement by creating a second touchpoint when customers return to redeem them. Fong (2017) reports that precisely tailored email promotions, while lifting focal product sales, can inadvertently reduce customers' search and cross-category exploration on the retailer's site, a possible unintended consequence of narrow targeting. In contrast, broader promotions or loyalty rewards may generate spillover purchases across categories. Our field experiment contributes to this literature by evaluating not just a one-time coupon incentive, but the optimal frequency of ongoing promotional emails. We provide experimental evidence on how reducing the cadence of coupon emails impacts purchase behavior and customer retention. Unlike prior coupon studies focusing mainly on immediate redemption and sales lift (e.g., Reimers and Xie 2019), we measure downstream outcomes (unsubscribe rates and repeat purchasing) to assess long-term efficacy. By embedding our frequency experiment in a live operations/marketing setting, our study complements earlier field experiments and demonstrates that controlled tests of policy changes (here, contact frequency) can reveal dynamic customer responses beyond the first transaction. This offers a bridge between short-run experimentation and long-run customer management outcomes.

2.3. Promotion/Coupon Optimization

Researchers in both operations and marketing have long sought to determine optimal promotion strategies. On the analytical side, models of pricing and couponing have addressed questions like which customers to target, when to send offers, and how deep discounts should be, all aiming to maximize profit or customer value under various constraints. For instance, Shaffer and Zhang (2002) analyze competitive one-to-one promotions, showing how firms can tailor coupon values to individual consumers in equilibrium. In a similar vein, Besanko et al. (2003) study price discrimination via promotions in a supply channel, finding that targeted discounts can be profitable even with aggregate data if designed correctly. Our work contributes a novel integration of empirical response estimation and optimization. We first use experimental and historical data to estimate how promotion frequency affects both immediate spending and churn risk, then embed these estimates in a customer lifetime value (CLV) optimization framework. This approach builds on recent advancements in data-driven promotion optimization (e.g., Jagabathula et al. 2022, Liu 2023, Chitla et al. 2023), which leverage machine learning and optimization to personalize offers. In particular, Jagabathula et al. (2022) propose a directed graph model to capture individual preferences for retail promotions and solve for personalized discount strategies, and Liu (2023) demonstrates a reinforcement learning method to dynamically target coupons in an online retail context. Also, our findings complement recent practice-based research (e.g., Baardman et al. 2019, Cohen et al. 2021, Baardman et al. 2023, Lin et al. 2024), where the authors develop promotion-focused algorithms and managerial frameworks that can be used in practice to optimally time promotional offers and to determine the value of both customized and mass promotions. We extend this stream by focusing on email contact frequency as the decision lever and by incorporating retention explicitly into the objective. Methodologically, our contribution lies in showing how causal impact estimates (from our field trial) can inform optimization to yield implementable policies.

2.4. Customer Lifetime Value and Personalization

Our research also contributes to the broad literature on customer lifetime value (CLV) and personalized marketing. It is well established that maximizing CLV often requires balancing short-term revenue against investments in customer retention and development (Reinartz and Kumar 2003). Firms have increasingly turned to personalization in marketing communications as a means to improve CLV by tailoring offers and contact strategies to individual preferences (Murthi and Sarkar 2003). Prior studies have explored various dimensions of personalization, from pricing to product recommendations, generally finding that personalization can enhance engagement and value if done intelligently. For example, Fader et al. (2005) introduce models linking individual purchase history (recency, frequency, monetary value) to CLV, enabling more effective targeting of marketing efforts. Venkatesan and Kumar (2004) empirically show that firms allocating marketing resources based on CLV metrics (rather than purely current profitability) achieve higher long-run profitability, underscoring the importance of forward-looking customer management. Moreover, personalization of promotional contacts has been examined in digital contexts: Sahni et al. (2017) and Todri et al.

(2020) both note that tailoring and frequency decisions should account for a customer's stage in the lifecycle or funnel position to avoid negative responses. However, operationalizing these insights remains challenging. Our study makes a practical contribution by demonstrating a data-driven approach to personalization of contact frequency. We directly estimate how each customer segment responds to more or fewer emails and use those heterogeneous effects to personalize an optimal contact policy. By comparing a myopic revenue-maximizing email strategy to a CLV-maximizing strategy, we quantify the uplift in lifetime value from adopting a retention-aware, personalized approach. In essence, we show that integrating individual-level churn risk into promotion decisions can meaningfully increase CLV, beyond what traditional one-size-fits-all or short-term optimized policies achieve. This extends the CLV literature by moving from descriptive models of value (e.g., Reinartz and Kumar 2003) to prescriptive analytics for value maximization. In summary, our research bridges multiple streams by using personalization (customized email frequencies) as a lever to improve CLV, supported by experimentation and optimization. We thus provide evidence that a rigorous CLV-based personalization of email frequency can simultaneously mitigate customer annoyance and enhance long-term profitability, contributing new insights to the intersection of customer analytics and operational marketing strategy.

3. Data and Empirical Setting

We collaborate with a major North American online retailer to study how coupon email engagement and annoyance jointly shape customer behavior. As a first step, we conducted a large-scale randomized field experiment in late October 2024, in which 603,153 customers were randomly assigned over a one-week period (October 25 to October 31, 2024) to one of two email-frequency conditions. Customers in the control group received daily coupon promotional emails, while customers in the treatment group received one email every two days. This randomized design provides clean causal evidence on how reducing email frequency affects customer engagement, unsubscriptions, and purchasing outcomes. Second, we analyzed a rich historical dataset tracking every coupon email and purchase for roughly 200,000 customers over the full year of 2024, allowing us to dive deeper into longer-run patterns and dynamics.

Across both datasets, we observe three core elements: email frequency, customer unsubscriptions, and purchasing behavior. Specifically, we track every coupon promotional email sent to each customer, how customers respond (for example, whether they open the email and whether they unsubscribe), and transaction-level purchase outcomes, including revenue, item quantity, and coupon usage. Taken together, these measures allow us to directly connect the platform's email-frequency strategy to unsubscription behavior and, ultimately, to long-run revenue outcomes.

3.1. Descriptive Statistics and Model-Free Evidence

Table 1 summarizes the historical dataset. Customers receive, on average, 0.94 coupon emails per day and open 23% of the emails they receive. The quality of coupons is relatively stable across customers, with an average coupon value of 18.3% with a standard deviation of 1.3%. Over the year, customers spend $269.86 on average and purchase 2.17 items, with an average per-item price of $172.60. The customer base is relatively mature, with an average tenure of 4.53 years, but it also includes many new users: 27% first visited in 2024 and an additional 16% first visited in 2023. Unsubscriptions are common: 15.6% of customers receiving emails unsubscribed at some point during the year, after which they stopped receiving coupon emails unless they subsequently resubscribed.

Table 1 Descriptive Statistics for Historical Dataset Covering One Year in 2024

Number of customers	200,000	
Customer-level statistics		
	Mean	SD
Emails Per Day	0.94	0.73
Email Open Rate	0.232	0.239
Coupon Value	0.183	0.013
Revenue ($/Year)	269.86	585.82
Item Quantity	2.17	3.45
Basket Size (Items/Order)	1.59	1.23
Price Per Item	172.60	257.35
Tenure (Years on Platform)	4.53	5.39
Unsubscribed	0.156	0.363

We next present model-free patterns from both datasets. Using the historical data, Figure 1 shows a sharp decline in purchasing after customers unsubscribe: average daily revenue is 59% lower than for customers who never unsubscribe, and 75% lower than the same customers' own spending prior to unsubscribing. Turning to the field experiment, Figure 2 plots unsubscription rates, revenue, and purchase rates in the weeks before, during, and after the experiment. Outcomes are similar in the pre- and post-periods, but diverge during the experimental week. Relative to daily emails, the reduced-frequency treatment decreases the unsubscription rate by 59.2% and reduces average daily revenue by 7.6% during the experimental week.

Taken together, these descriptive patterns highlight a central tradeoff: higher email frequency increases short-run revenue but also raises unsubscriptions, which are associated with a large

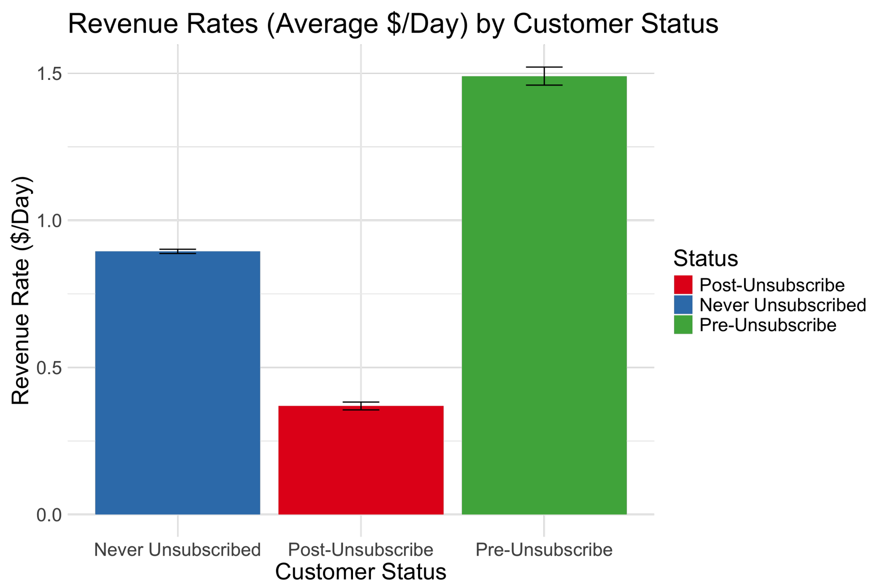

Figure 1 **Revenue rates significantly decrease after customers unsubscribe.**

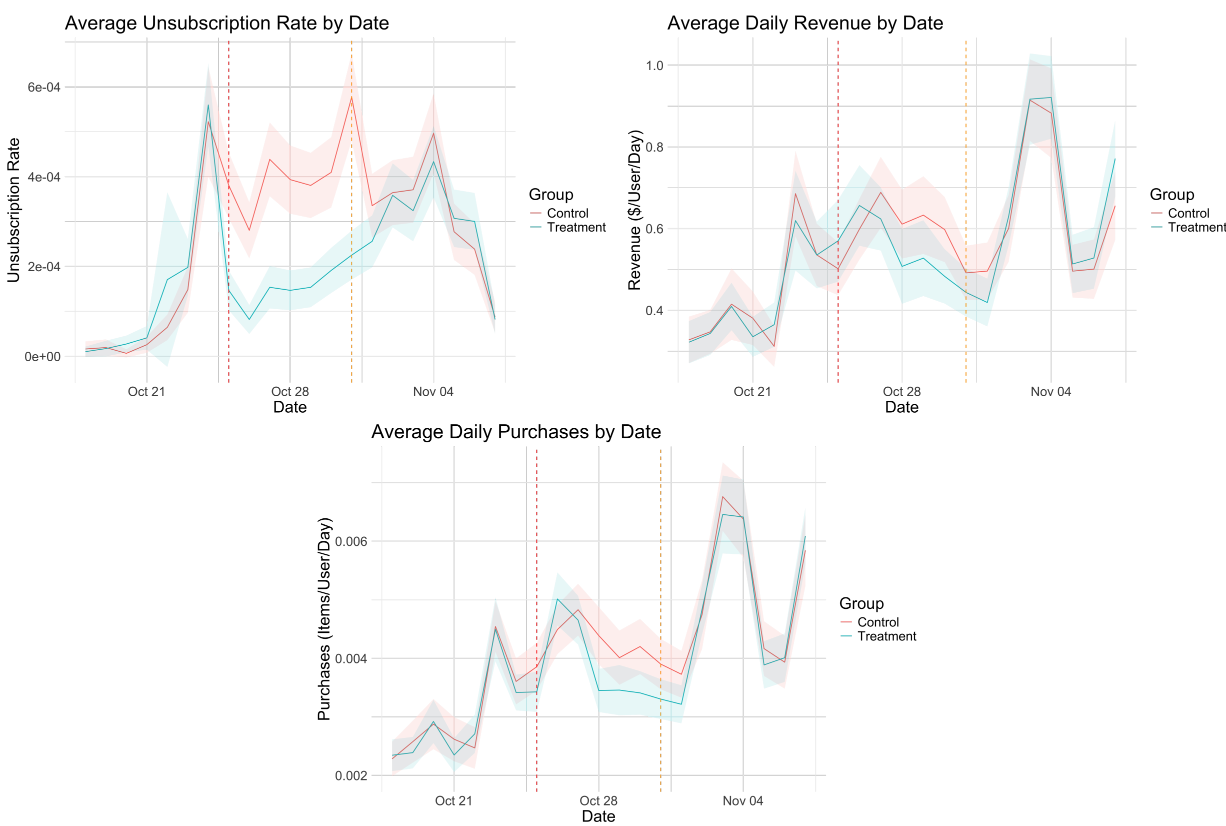

Figure 2 **Sending fewer emails significantly decreases unsubscription rates and slightly decreases short-term revenue and purchase rates. The red dashed line marks the start of the field experiment, and the orange line marks the end of the experiment. Shaded areas indicate 95% confidence intervals for each date.**

decline in subsequent purchasing. Because customers' email exposure, unsubscription decisions, and spending may be jointly determined, we rely on the randomized field experiment to identify the causal impact of email frequency on unsubscriptions and short-term purchasing outcomes. We then complement the experimental evidence with the historical data to characterize the longer-run relationships between unsubscriptions, customer retention, and long-term revenue.

4. Results

We present findings from both a field experiment and historical data analysis to examine the relationships between email frequency, customer unsubscriptions, and long-term revenue. First, the field experiment reveals a tradeoff: a higher email frequency boosts short-term revenue but also leads to higher unsubscription rates, while reducing email frequency has the opposite effect. Second, our analysis of historical data indicates that unsubscriptions lead to decreased revenue and lower customer retention, ultimately negatively impacting long-term customer value.

4.1. Field Experiment: Main Effects on Engagement and Sales

We begin with the core results of the field experiment, which adjusted the frequency of coupon emails. Given the random assignment of customers[2], our baseline specification compares outcomes between the treated group, which received emails every other day, and the control group, which received daily emails:

$$\text{Outcome}_{it} = \text{Treated}_i \mid \text{Date}_t \tag{1}$$

Our analysis focuses on two classes of outcomes: (i) engagement and churn outcomes reflected by customer unsubscribing from the email list, which we interpret as a revealed measure of annoyance and an early signal of attrition, and (ii) short-run sales outcomes (purchase incidence, number of items purchased, and revenue generated over the experimental horizon). Table 2 reports the average treatment effects of reducing email frequency from daily to every-other-day on these outcomes during the one-week experiment. Overall, the results reveal a clear retention-revenue trade-off:

- *Unsubscription rate.* Customers assigned to the every-other-day email frequency were significantly less likely to unsubscribe during the experiment week than customers who received daily emails. The estimated effect corresponds to a *59.2% relative reduction* in the probability of unsubscribing.[3] From a managerial perspective, this provides causal evidence that aggressive messaging can push some customers away, whereas holding back communications can keep more customers engaged (at least in the sense of remaining subscribed).

[2] Note that we verified the randomness of the treatment assignment in Appendix C.

[3] In absolute terms, the baseline (daily) unsubscribe rate is low, on the order of a few tenths of a percent per week, reflecting that our experiment spans only a short horizon (one week). Even so, halving email frequency more than halves this unsubscription probability, and these seemingly modest weekly differences can compound into meaningful retention gains over longer time horizons. Over one year, the unsubscription rate drops from 12.6% to 5.1%.

- *Short-run purchase incidence and revenue.* The reduction in email frequency comes at a cost to short-run sales. Relative to the daily-email control group, treated customers (every other day) purchased less often and generated lower revenue during the experiment week. Specifically, the probability of making a purchase on a given day declines by 4.8% relative to the daily-email group, and daily revenue per customer falls by 7.6% on average under the reduced-frequency treatment. These declines imply that daily coupon emails stimulate incremental purchases in the immediate term. Intuitively, fewer emails imply fewer reminders and fewer opportunities to click through and shop, leading to a modest reduction in weekly sales among treated customers.
- *Price and number of items purchased.* Consistent with the aforementioned interpretation, we also observe a 10% decrease in the number of items purchased. By contrast, the average price per item is essentially unchanged across groups. Thus, the revenue decline is driven primarily by fewer purchase occasions (and slightly smaller baskets), rather than shifts in the composition or pricing of items purchased. In summary, sending emails daily generates slightly more frequent purchases and higher daily spending per customer than sending emails every other day, confirming the short-run revenue benefit of a higher contact frequency.

Taken together, these experimental results confirm a fundamental tension: emailing customers daily yields higher immediate revenue but induces significantly more churn from the email list; conversely, emailing less often protects against churn but sacrifices some short-run sales. For the retailer, this quantifies the engagement–annoyance trade-off: more aggressive outreach buys short-term engagement at the expense of longer-run relationship potential. Importantly, the magnitudes in our setting suggest that neither objective can be omitted: a nearly 60% reduction in churn is highly significant, while a 5–10% decrease in weekly sales is also salient to managers. The firm must therefore actively balance these competing outcomes. The results underscore that there is no free lunch: one cannot reduce email volume without some short-run revenue loss, but one also cannot push frequency to the maximum without eroding customers' willingness to engage over time.

Firms should recognize that overly frequent customer communications can boost short-run sales metrics while also increasing customer fatigue and churn (unsubscription). Our experiment provides concrete evidence of this trade-off: within a single week, halving email frequency substantially reduces opt-outs, preserving a larger pool of reachable customers for future campaigns, at the cost of a single-digit percentage drop in that week's revenue. Depending on a firm's planning horizon and objectives, this trade-off may be worthwhile, especially when long-run retention is a priority. The managerial challenge is to identify a sweet spot in contact frequency that sustains engagement while minimizing burnout. In the next sections, we examine heterogeneity in these effects to clarify which customers are most sensitive to email frequency and how firms can tailor outreach accordingly.

Table 2 Average Treatment Effects for Reducing Email Frequency

	Dependent variable:				
	P(Unsubscribe)	Revenue ($/Day)	Item Quantity	P(Purchase)	Price Per Item
	(1)	(2)	(3)	(4)	(5)
Treated	-0.0002^{***}	-0.037^{**}	-0.0004^{***}	-0.0001^{**}	0.671
	(0.00002)	(0.018)	(0.0001)	(0.00005)	(4.021)
Observations	4,222,071	4,222,071	4,222,071	4,222,071	9,635
Rel. Effect Size	-59.2%	-7.6%	-10.0%	-4.8%	+0.3%
Fixed Effects	Date	Date	Date	Date	Date
Note:					*p<0.1; **p<0.05; ***p<0.01

4.2. Field Experiment: Heterogeneous Effects by Customer Segment

Customers do not respond uniformly to frequent coupon emails. Therefore, one of the contributions of our study is to unpack heterogeneous treatment effects (HTEs) of the reduced-frequency policy, identifying which segments realize the largest retention gains (lower unsubscription) and which segments experience the largest short-run revenue losses when email frequency is reduced. We focus on three moderators motivated by managerial relevance and the underlying behavioral mechanisms: (i) customer tenure (relationship maturity and loyalty), (ii) recent purchase volume (proxied by item quantity purchased, measuring customer purchasing activity), and (iii) email engagement (historical open rate). These moderators capture distinct dimensions of heterogeneity: tenure distinguishes newer versus established customers; purchase volume separates high-activity customers from light shoppers; and open rate measures responsiveness to email content. For each moderator, we estimate interaction models that augment the baseline specification in Table 2 with a treatment-by-moderator term, and we corroborate the regression-based patterns using median-split subgroup analyses reported in Appendix D.

4.2.1. Heterogeneity by customer tenure. Customer tenure is an important moderator of the email-frequency effect. Table 3 shows that reducing email frequency yields substantially larger retention benefits for newer customers than for long-tenured customers. In interaction regressions, the treatment-by-tenure term for unsubscription is positive and statistically significant, indicating that the unsubscribe-reducing effect of the treatment attenuates with tenure. Put differently, for a customer with many years of tenure, moving to an every-other-day cadence yields only a modest improvement in retention, whereas for a newly acquired customer, the same reduction can meaningfully increase the likelihood of remaining subscribed. This finding is consistent with the following theory: newer customers are more fragile and more easily pushed away by aggressive contact, whereas long-tenured customers are more tolerant of frequent emails but also more likely to respond to them behaviorally. Managerially, this pattern suggests that firms should consider easing in new customers with a lighter touch to reduce early churn, while maintaining a higher frequency for established customers who have demonstrated tolerance and responsiveness.

Table 3 HTEs: Customer Tenure

	Dependent variable:	
	P(Unsubscribe)	Revenue (\$/Day)
	(1)	(2)
Tenure	-0.00000^{***}	0.00002^{**}
	(0.000)	(0.00001)
Treated	-0.0003^{***}	-0.077^{**}
	(0.00003)	(0.036)
Tenure×Treated	0.00000^{**}	0.00001
	(0.000)	(0.00001)
Observations	4,222,071	4,222,071
Note:		*p<0.1; **p<0.05; ***p<0.01

4.2.2. Heterogeneity by purchase volume. We next examine heterogeneity by recent purchase activity, using pre-experiment item quantity as a proxy for purchase volume. This moderator separates high-activity customers from casual shoppers. It follows from Table 4 that the trade-off differs sharply across these segments: high-volume customers exhibit little to no retention benefit from reduced frequency but experience substantial short-run revenue losses when emails are less frequent.

Customers with high recent purchase volume are generally unlikely to opt out even under daily emails, plausibly because they derive value from frequent coupons and have a strong ongoing interest in the retailer. The treatment effect on short-run revenue becomes substantially more negative as recent purchase volume increases. In the interaction models, the treatment-by-volume coefficient for revenue is strongly negative and statistically significant, implying that cutting email frequency results in spending more for heavy buyers than for light shoppers. Median-split results in Appendix D reinforce this pattern: revenue declines are concentrated among the most active customers, whereas low-activity customers exhibit only modest sales differences between daily and every-other-day policies.

This combination implies that heavy buyers are relatively insulated from email fatigue (they rarely unsubscribe), yet they are highly responsive to purchase triggers embedded in email content.

Light buyers, by contrast, contribute relatively little incremental revenue from additional emails but are more likely to disengage when over-contacted. From a managerial standpoint, these results imply that maintaining a higher frequency for heavy buyers can preserve revenue without substantially increasing churn, while moderating frequency for low-volume customers can deliver meaningful retention gains at limited short-run revenue cost.

Table 4 HTEs: Item Quantity in Past 7 Days

	Dependent variable:	
	P(Unsubscribe)	Revenue (\$/Day)
	(1)	(2)
Item Quantity	0.00003	6.632^{***}
	(0.00002)	(0.029)
Treated	-0.0002^{***}	0.093^{***}
	(0.00002)	(0.018)
Item Quantity×Treated	0.00002	-6.189^{***}
	(0.00005)	(0.057)
Observations	4,222,071	4,222,071
Note:		*p<0.1; **p<0.05; ***p<0.01

4.2.3. Heterogeneity by email open rate. Finally, we study heterogeneity by historical email open rate, which captures customer engagement with email content. We find a nuanced pattern in Table 5: highly engaged customers realize especially large retention benefits from reduced frequency and experience comparatively small revenue losses, whereas low-engagement customers exhibit a smaller change in unsubscription but can experience a larger proportional reduction in purchases when frequency is reduced.

Among high open-rate customers, the reduced-frequency policy substantially lowers unsubscription. The interaction between treatment and open rate for unsubscription is negative and statistically significant, indicating that the retention benefit of reducing email frequency is amplified for customers who actively read emails. This pattern suggests that even engaged customers can become fatigued when contact is too frequent; because they pay attention to email content, daily emails may feel repetitive or overly annoying, increasing the likelihood of opting out. A more moderate email strategy provides relief from inbox clutter and preserves goodwill among these attentive customers.

On the revenue side, high open-rate customers do not exhibit large spending declines when frequency is reduced. In the interaction regressions, the treatment-by-open-rate term for revenue is positive, implying smaller revenue losses at higher engagement. Intuitively, customers who open most of what they receive continue to be exposed to offers even at an every-other-day email frequency, so the marginal reduction in purchase triggers is limited. Moreover, for this group, sending fewer emails may even increase the effectiveness of each email by avoiding oversaturation.

Low open-rate customers show the opposite structure. Because they rarely engage with emails, unsubscription is less strongly affected by reduced email frequency: these customers often ignore messages rather than formally opting out, so reducing frequency yields little retention benefit. Yet fewer emails also mean fewer opportunities to catch their attention on the occasional day they might open and purchase, leading to a larger proportional decrease in conversions for this segment.

Overall, the open-rate results underscore that the retention–revenue trade-off depends not only on customer value but also on attention and responsiveness. Reducing frequency appears particularly attractive for highly engaged customers, delivering sizable retention gains with limited short-run revenue impact.

Table 5 HTEs: Email Opens in Past 7 Days

	Dependent variable:	
	P(Unsubscribe)	Revenue ($/Day)
	(1)	(2)
Email Opens	0.0001^{***}	-0.005
	(0.00001)	(0.010)
Treated	-0.0001^{***}	-0.062^{***}
	(0.00002)	(0.022)
Email Opens×Treated	-0.0001^{***}	0.026^{*}
	(0.00001)	(0.014)
Observations	4,222,071	4,222,071
Note:		*p<0.1; **p<0.05; ***p<0.01

Analyzing the HTEs in conjunction, Table 6 shows similar results to our individual HTE analyses, with customer tenure and email opening frequency as the significant effects on unsubscription rates, and recent purchasing behavior as the dominant effect on revenue.

4.2.4. Managerial implications. The heterogeneous effects documented in this section reinforce a central conclusion: the optimal email policy is not one-size-fits-all. High-value, long-tenured, and high-volume customers tend to be more revenue-responsive and less churn-sensitive, whereas new and low-volume customers exhibit stronger retention gains from moderation. Email engagement further refines this segmentation by identifying customers for whom high-frequency contact

Table 6 HTEs: All Variables Combined

	Dependent variable:	
	P(Unsubscribe)	Revenue (\$/Day)
	(1)	(2)
Tenure	-0.00000^{***}	0.00002^{**}
	(0.000)	(0.00001)
Item Quantity	0.00003	6.632^{***}
	(0.00003)	(0.029)
Email Opens	0.0001^{***}	0.007
	(0.00001)	(0.010)
Treated	-0.0002^{***}	0.042
	(0.00003)	(0.038)
Tenure×Treated	0.00000^{***}	0.00001
	(0.000)	(0.00001)
Item Quantity×Treated	0.00002	-6.189^{***}
	(0.00005)	(0.057)
Email Opens×Treated	-0.0001^{***}	0.015
	(0.00001)	(0.014)
Observations	4,222,071	4,222,071
Note:		*p<0.1; **p<0.05; ***p<0.01

risks fatigue without significant revenue benefit. We leverage these patterns in Section 5 when constructing and evaluating segment-based (personalized) email policies.

4.3. Historical Data Analyses: Effects of Unsubscribing on Long-Term Behavior

In the historical data, we track customer behavior over a one-year period in 2024 to assess the longer-run consequences of unsubscriptions. We aggregate outcomes to the customer-month level and focus on the event of a customer opting out of the retailer's promotional emails. Unsubscribing is an explicit disengagement decision that both signals fatigue and mechanically removes a key marketing touchpoint. Our central question is whether, and by how much, this opt-out changes subsequent purchasing and retention—and quantifying these effects is essential for evaluating the long-run cost of driving customers to the point of disengagement.

We next examine the longer-run consequences of customer annoyance and disengagement using the historical panel. Our focal event is a customer unsubscribing from the retailer's promotional emails. Unsubscribing represents an explicit opt-out decision that both reflects disengagement and mechanically removes a key marketing touchpoint. The central question is whether, and by how much, this event changes subsequent purchasing and retention. Quantifying these effects is essential for assessing the long-run cost of pushing customers to the point of disengagement. Equations (2)-(4) show our specifications estimating the effects of unsubscribing on revenue and survival with observations at the customer-month level.

$$\text{Revenue}_{it} = \text{Unsubscribed}_{it-1} + \text{Revenue}_{it-1} \mid \text{Month}_{t-1} \tag{2}$$

$$\text{Revenue}_{it} = \text{Unsubscribed}_{it-1} + \text{Revenue}_{it-1} \mid \text{Month}_{t-1} \mid \text{EmailFrequency}_{it-1} \tag{3}$$

$$\text{Survival}_{it} = \text{Unsubscribed}_{it-1} + \text{Tenure}_{it-1} + \text{Unsubscribed}_{it-1} \times \text{Tenure}_{it-1} \tag{4}$$

Using the panel data, we compare purchase trajectories before versus after the unsubscribe event and contrast customers who unsubscribe with otherwise similar customers who remain subscribed, controlling for rich observables and prior behavior. Table 7 summarizes the core results for monthly revenue, and Table 8 reports corresponding effects on customer retention. Because customers who unsubscribe may differ systematically from those who do not, we complement descriptive before-and-after patterns with regression specifications that include controls for pre-unsubscribe purchase levels in (2). We further implement an instrumental-variables strategy to mitigate selection concerns in (3); specifically, we leverage the email frequency variation as instruments for unsubscribe behavior in a holdout period. The instrument satisfies the relevance condition because it affects the probability of unsubscribing in the previous month, and satisfies the exclusion restriction because email frequency in the previous month does not directly affect revenue in the current month[4] (see Appendix E for robustness and details, including first-stage evidence of instrument relevance). Finally, we estimate the effects of unsubscribing on customer retention after controlling for customer tenure in (4).

Across specifications, the results are consistent: unsubscribing is associated with a large and statistically significant decline in future spending and engagement. Economically, the magnitude is substantial. In the month after unsubscribing, expected customer revenue drops by roughly one-third to one-half relative to the counterfactual in which the same customer remained subscribed. For instance, a representative baseline customer spends an average of \$22.47 per month; after unsubscribing, spending declines by \$9.85 on average. This corresponds to a 44% reduction in a simple before-after comparison. Our preferred IV estimates attribute a 36% decrease in monthly revenue to unsubscribing, consistent with the magnitude of the OLS regression. Put differently, if a customer would have spent \$100 over the coming months, unsubscribing reduces expected spending

[4] Note that the coupon codes distributed via email are valid for a day or two. As a result, coupons sent in the previous month cannot directly affect customer spending in the current month.

to about \$66 on average. Figure 1 provides complementary visual evidence, showing a sharp drop in daily spending immediately after an unsubscribe event; the regression and IV analyses confirm that this pattern is not driven solely by spurious correlation.

Table 7 Unsubscribing Significantly Decreases Revenue

	Dependent variable:	
	Month Revenue (\$)	
	(1)	(2)
Previous Month Revenue	0.086***	0.086***
	(0.001)	(0.001)
Previous Month Unsubscribed	−9.853***	−8.097***
	(0.745)	(0.980)
Observations	1,095,832	1,095,832
Fixed Effects	Month	Month
IV	None	Previous Month Email Frequency
Note:		*p<0.1; **p<0.05; ***p<0.01

Unsubscribing also strongly impacts reduced retention in an active purchasing state. We define a customer as retained in a given time period, or surviving, if they make at least one purchase during that period. Relative to customers who remain subscribed, customers who unsubscribe exhibit a markedly lower probability of making any purchase in subsequent months. Using panel regressions for survival (Table 8), we estimate that the likelihood of being active over the six months after unsubscribing is more than 60% lower than for otherwise similar customers who remain subscribed, with the gap persisting across months. The first column uses retention in months 2-12 for customers who unsubscribe in month 1, while the second column uses retention in months 7-12 for customers who unsubscribe in months 1-6. In practical terms, once customers opt out of emails, a large share do not return to transact in the near future, making unsubscribe a strong leading indicator of downstream churn from the retailer.

Overall, the evidence implies a sizable long-run revenue penalty when a customer disengages to the point of unsubscribing. This historical analysis therefore provides the critical link between the short-run experimental trade-off documented in Table 2 and long-run customer value: opt-outs are not merely a transitory annoyance outcome, but a precursor to economically meaningful spending and retention losses.

We emphasize that unsubscribing can be both a cause and a symptom. Customers who unsubscribe may already be on a downward trajectory, and the act of unsubscribing simultaneously

Table 8 Effect of Unsubscribing on Customer Retention

	Dependent variable:	
	Survival	
	After Month 1	After Months 1-6
	(1)	(2)
Unsubscribed	-0.116^{***}	-0.051^{***}
	(0.030)	(0.004)
Tenure	0.00003^{***}	0.00002^{***}
	(0.00000)	(0.00000)
Unsubscribed×Tenure	-0.00002	-0.00001^{***}
	(0.00002)	(0.00000)
Constant	0.183^{***}	0.084^{***}
	(0.004)	(0.001)
Observations	23,465	118,941
Note:		*p<0.1; **p<0.05; ***p<0.01

removes future marketing triggers that could have supported purchasing. Our fixed-effects and trend controls, together with the IV approach, indicate that a large decline remains even after adjusting for selection, supporting a substantial causal component. From a managerial perspective, unsubscribe is therefore a high-salience warning signal of reduced future value.

4.3.1. Heterogeneous effects of unsubscribing. The long-run consequences of unsubscribing may also differ across customer segments. Using the same moderators as in Section 4.2 (tenure, purchase volume, and prior email engagement), we assess whether the post-unsubscribe declines in revenue and retention vary systematically across customers. The results are presented in Table 9. We do not find statistically significant differences by tenure or by prior email engagement. In contrast, purchase volume meaningfully moderates the unsubscribe effect. This pattern is consistent with a positive interaction between prior purchase volume and unsubscribing, suggesting that higher-value customers are less likely to fully disengage even after opting out of email communications.

These heterogeneous patterns underscore the value of targeted retention strategies. When highly engaged customers unsubscribe, the firm loses a particularly effective communication channel, and the subsequent spending decline can be substantial. This makes it important to mitigate fatigue among these customers through more careful management of digital outreach. Taken together with the field experiment results in Table 2, the historical analyses point to a consistent mechanism: higher email frequency boosts short-run sales but also increases opt-outs, and opt-outs are followed by economically large and persistent declines in spending and active retention. In the next section, we formalize this short-run versus long-run trade-off and use it to develop a customer-segmented framework for optimizing email frequency to maximize customer lifetime value (CLV).

Table 9 HTEs for Effect of Unsubscribing on Revenue

	Dependent variable:		
	Month Revenue ($)		
	(1)	(2)	(3)
Previous Month Revenue	0.087***	0.087***	0.060***
	(0.001)	(0.001)	(0.001)
HTE	0.001***	0.052***	7.378***
	(0.0001)	(0.004)	(0.148)
Previous Month Unsubscribed	−6.597***	−10.222***	−8.465***
	(1.281)	(1.313)	(1.024)
Previous Month Unsubscribed×HTE	−0.001	0.035	4.861***
	(0.001)	(0.036)	(0.968)
Observations	1,095,832	1,095,832	1,095,832
HTE	Tenure	Email Opens	Item Quantity
Fixed Effects	Month	Month	Month
IV	Yes	Yes	Yes
Note:			*p<0.1; **p<0.05; ***p<0.01

4.4. Robustness Checks

We perform a series of robustness checks for our main results, showing robustness to alternate variable definitions and model specifications. We summarize the robustness results here and show the full results in the appendices.

Appendix A shows that the results for the field experiment and historical data analyses are robust to log-transforming key dependent variables. Appendix B shows that the results are robust to alternative model specifications such as logistic and Poisson regressions. Appendix C shows that the field experiment treatment group is randomly assigned and the results are robust to a difference-in-differences approach. Appendix D shows that the HTE results are robust to binary transforming customer characteristics on a median split. Appendix E shows that the historical data IV of past month email frequency is relevant and robust to alternative IV specifications.

5. Optimization of Email Frequency for Customer Lifetime Value

Having established the empirical relationships between email frequency, short-term revenue, and long-term customer retention, we next integrate these components into a customer lifetime value optimization model. The goal of this optimization is to determine an optimal frequency for the email policy that maximizes the long-term objective of a company, which is customer lifetime value, leveraging the causal estimates we obtained. In essence, we want to answer: How should the retailer personalize or adjust its email frequency across different customers to maximize total customer lifetime value, and how much improvement can this yield relative to the current policy of emailing everyone daily?

5.1. Model, Objective, and Decision Variables

Our goal is to maximize long-run profitability by personalizing email frequency at the customer level. The platform chooses, for each customer i, whether to send coupon emails daily or every other day. Increasing email frequency raises short-run purchase probability and revenue, but it also increases the likelihood of unsubscribing, which reduces future revenue by shortening the customer's effective lifetime. Moreover, for a subset of customers, oversaturation can reduce responsiveness, so a lower frequency may improve even short-run revenue.

Formally, let $Y_i \in \{0,1\}$ denote the email-frequency assignment for customer i, where $Y_i = 0$ corresponds to every-other-day emails and $Y_i = 1$ corresponds to daily emails. Let X_i collect customer covariates (tenure, email open rate, and past-month item quantity). We consider a constrained optimization that limits total coupon usage to a budget B:

$$\begin{aligned} \max_{\{Y_i\}_{i=1}^N : Y_i \in \{0,1\}} \quad & \sum_{i=1}^N CLV(X_i, Y_i) \\ \text{s.t.} \quad & \sum_{i=1}^N \mathbb{E}[\text{CouponUsage} \mid X_i, Y_i] \le B. \end{aligned} \tag{5}$$

We also study the unconstrained variant of (5) by dropping the budget constraint.

5.2. CLV Construction and Identification of Model Components

We define customer lifetime value as the steady-state expected value of future revenue under an email policy, accounting for churn risk:

$$CLV(X,Y) = \frac{\mathbb{E}[\text{Revenue} \mid X, Y]}{1 - \mathbb{P}(\text{Survival} \mid X, Y) + r}, \tag{6}$$

$$\mathbb{E}[\text{Revenue} \mid X, Y] = \mathbb{E}[\text{Revenue} \mid U, X, Y] \cdot \mathbb{P}(U \mid X, Y) \; + \; \mathbb{E}[\text{Revenue} \mid NU, X, Y] \cdot \big(1 - \mathbb{P}(U \mid X, Y)\big),$$

$$\mathbb{P}(\text{Survival} \mid X, Y) = \mathbb{P}(\text{Survival} \mid U, X) \cdot \mathbb{P}(U \mid X, Y) \; + \; \mathbb{P}(\text{Survival} \mid NU, X) \cdot \big(1 - \mathbb{P}(U \mid X, Y)\big),$$

where U indicates that the customer unsubscribes (email churn) and NU indicates that the customer remains subscribed. Throughout, we set the discount rate to $r = 0$, which is appropriate given the relatively short horizon over which survival probabilities become small in our application.

We estimate each component of (6) by combining experimental and historical evidence. From the field experiment, we estimate the causal effect of frequency Y on (i) revenue among subscribed customers, $\mathbb{E}[\text{Revenue} \mid NU, X, Y]$, (ii) unsubscription risk, $\mathbb{P}(U \mid X, Y)$, and (iii) coupon usage, $\mathbb{E}[\text{CouponUsage} \mid X, Y]$. From historical data, we estimate the longer-run consequences of unsubscribing for (i) post-unsubscribe revenue and (ii) survival dynamics, summarized by $\mathbb{E}[\text{RevenueDecrease} \mid U]$, $\mathbb{P}(\text{Survival} \mid U, X)$, and $\mathbb{P}(\text{Survival} \mid NU, X)$. Finally, we impute revenue for unsubscribed customers as

$$\mathbb{E}[\text{Revenue} \mid U, X, Y] = \mathbb{E}[\text{Revenue} \mid NU, X, Y] \times \mathbb{E}[\text{RevenueDecrease} \mid U]. \tag{7}$$

5.3. Policy Classes and Optimization Variants

To clarify the role of retention in the policy design, we compare two optimization approaches.

Myopic (Short-Run) Optimization. The myopic approach selects Y_i to maximize the short-run revenue rate $\mathbb{E}[\text{Revenue} \mid NU, X_i, Y_i]$ while ignoring churn effects. This captures a revenue-first policy that can still personalize frequency when some customers respond better to lower contact intensity.

CLV-Optimal (Retention-Aware) Optimization. The full approach selects Y_i to maximize $CLV(X_i, Y_i)$ as defined in (6), explicitly trading off short-run revenue against churn risk.

For each approach, we consider: (i) an unconstrained version (no budget on coupon usage), and (ii) a constrained version with B set to match projected status-quo coupon usage.

5.4. Results: Unconstrained and Constrained Policies

Table 10 Unconstrained Optimization Results

Email Policy	CLV ($)	Revenue Rate ($/Month)	E[Revenue] ($/Year)	Survival (/Year)	Coupon Usage ($/Year)	Email Rate (/Day)
Every Day (Current)	213.03	11.78	130.3	0.376	12.42	1
Every Other Day	213.56	10.94	127.2	0.388	11.55	0.50
Myopic Optimization	227.46	12.12	136.9	0.382	17.61	0.68
Full Optimization	230.78	12.04	138.2	0.388	19.85	0.52

Table 10 reports unconstrained results. Under the status quo (daily emails for all), average CLV is \$213.03. A uniform switch to every-other-day emails yields a slightly higher CLV (\$213.56; +0.2%), reflecting modest retention gains that approximately offset the short-run revenue loss. Importantly, switching to every-other-day emails delivers a slightly higher CLV while reducing coupon usage by as much as 7%.

Personalization delivers substantially larger gains. The myopic policy increases CLV to $227.46 (+6.8% vs. status quo), with improvements driven by higher annual revenue (+5.1%) and a modest increase in survival (+1.6%). Notably, the myopic policy can assign a lower frequency to customers who are more responsive under a less frequent email policy, consistent with oversaturation effects. These gains come with higher expected coupon usage ($17.61 vs. $12.42 per customer-year), reflecting that targeting responsiveness increases promotional activity.

The CLV-optimal policy further increases CLV to $230.78 (+8.3% vs. status quo; +1.4% vs. myopic). Relative to the myopic policy, the retention-aware email assignment leads to a slightly lower short-run revenue rate but achieves higher long-run revenue by lowering unsubscribe risk and improving survival. The resulting policy also features a lower average email rate than the myopic policy (0.52 vs. 0.68 per day), illustrating that the model assigns reduced frequency where it most improves lifetime value, not merely where it reduces cost.

Table 11 Constrained Optimization Results

Email Policy	CLV ($)	Revenue Rate ($/Month)	E[Revenue] ($/Year)	Survival (/Year)	Coupon Usage ($/Year)	Email Rate (/Day)
Myopic Optimization	216.76	11.78	131.3	0.377	12.42	0.88
Full Optimization	222.53	11.65	133.1	0.384	12.42	0.63

Table 11 reports constrained results, where expected coupon usage is capped at the status-quo level ($12.42 per customer-year). Under this constraint, the myopic policy yields a modest CLV gain (+1.8%), whereas the CLV-optimal policy delivers a larger improvement (+4.5%). In both cases, the CLV gains reflect a combination of improved expected annual revenue and increased survival, but the retention-aware policy produces systematically larger gains because it allocates the fixed coupon budget toward customers for whom the incremental churn risk is most costly.

5.5. Heterogeneity and Segmentation Implications

To characterize which customers should receive reduced frequency, we compare daily versus every-other-day policies across median splits of tenure (new vs. old), past-year spending (high vs. low), and past-year promotion usage (high vs. low deal seekers). For spending and promotion usage, where the overall median is zero, we split at the median among customers with strictly positive values. Table 12 summarizes the resulting heterogeneity patterns.

For older customers, daily emails can dominate: the incremental revenue gain is comparatively large, while the retention penalty is comparatively small, yielding higher CLV under the daily email policy. For newer customers, reduced frequency tends to dominate: revenue losses are small and retention gains are larger, producing higher CLV under every-other-day email policy. Similarly, daily emails tend to be more attractive for high spenders and high coupon users, whereas reduced frequency performs better for low spenders and low coupon users, consistent with stronger churn sensitivity among less-established customers.

Table 12 Heterogeneous Treatment Effects for Optimization Results

Email Policy	CLV ($)	Revenue Rate ($/Month)	E[Revenue] ($/Year)	Survival (/Year)	Coupon Usage ($/Year)	Email Rate (/Day)
Old Customers (Age $>$ 2583 Days)						
Every Day (Current)	184.25	11.45	125.5	0.319	11.08	1
Every Other Day	178.90	10.27	119.3	0.332	9.38	0.50
New Customers (Age $\leq$ 2583 Days)						
Every Day (Current)	241.8	12.12	135.1	0.432	13.75	1
Every Other Day	248.3	11.60	135.08	0.445	13.71	0.50
High Spenders (Past Year Spending $>$ $179.49)						
Every Day (Current)	214.42	11.94	131.31	0.375	17.41	1
Every Other Day	213.28	10.94	127.13	0.388	12.48	0.50
Low Spenders (Past Year Spending $\leq$ $179.49)						
Every Day (Current)	211.79	11.65	129.3	0.376	7.94	1
Every Other Day	213.82	10.93	127.3	0.388	10.71	0.50
High Coupon Users (Past Year Coupon Usage $>$ $ 17.09)						
Every Day (Current)	221.11	12.35	134.70	0.379	20.26	1
Every Other Day	217.4	11.11	128.63	0.392	13.22	0.50
Low Coupon Users (Past Year Coupon Usage $\leq$ $ 17.09)						
Every Day (Current)	210.86	11.63	129.1	0.375	10.31	1
Every Other Day	212.53	10.89	126.8	0.387	11.10	0.50

5.6. Sensitivity to Baseline Retention

Finally, we examine how the value of reduced frequency and personalization varies with the platform's baseline customer retention. We scale the average observed annual survival rate (0.376) to values between 0.1 and 0.9, scaling the every-other-day survival rate proportionally. For example, when baseline survival is set to 0.9, the corresponding every-other-day survival is $0.9 \times \frac{0.388}{0.376} = 0.929$. To keep predicted CLV within a realistic range, we cap individual survival probabilities at 0.99.

Table 13 and Figure 3 show that reduced frequency becomes increasingly valuable as baseline retention rises. When customers are more loyal, protecting retention (and avoiding unsubscribe-driven churn) has a larger payoff in terms of lifetime value; conversely, when baseline churn is high, short-run revenue considerations become relatively more important. Consistent with this logic, the relative performance of every-other-day emails and of both optimization policies improves monotonically with baseline retention. Moreover, across the full range of baseline survival values, the CLV-optimal policy delivers the largest gains, highlighting the importance of explicitly accounting for unsubscribing when designing email-frequency policies.

Table 13 Ratio Between CLV for Optimization Strategies and Status Quo for Different Base Survival Values

Base Survival	Every Other Day	Myopic Optimization	Full Optimization
0.1	0.9811	1.0534	1.0648
0.2	0.9869	1.0571	1.0698
0.3	0.9946	1.0622	1.0764
0.376 (Actual)	1.0025	1.0677	1.0833
0.4	1.0056	1.0700	1.0860
0.5	1.0242	1.0843	1.1021
0.6	1.0868	1.1382	1.1562
0.7	1.2385	1.2760	1.2952
0.8	1.3712	1.3915	1.4194
0.9	1.5010	1.4928	1.5423

5.7. Discussion and Managerial Implications

Our findings carry several actionable implications for managers overseeing customer contact strategies and promotional email campaigns. First and foremost, we provide empirical evidence that more communication is not always better as there is a real cost to customer over-engagement, in the form of higher opt-out rates and lost future revenue. Managers should therefore resist the reflex to maximize short-term metrics (like campaign revenue or click-through rates) by bombarding customers with messages. Instead, firms should aim to adopt a more nuanced, retention-aware approach. Concretely, this could mean setting frequency caps or enforcing rules for email sends. For example, a retailer might implement a policy to limit mass promotional emails to every other day for certain customer segments, especially those who have shown signs of fatigue (such as purchasing less frequently over time or being relatively new to the list). Our results suggest that such a policy would reduce unsubscribes and extend customer lifetimes with minimal impact on immediate sales, which, in many cases, might improve overall profitability.

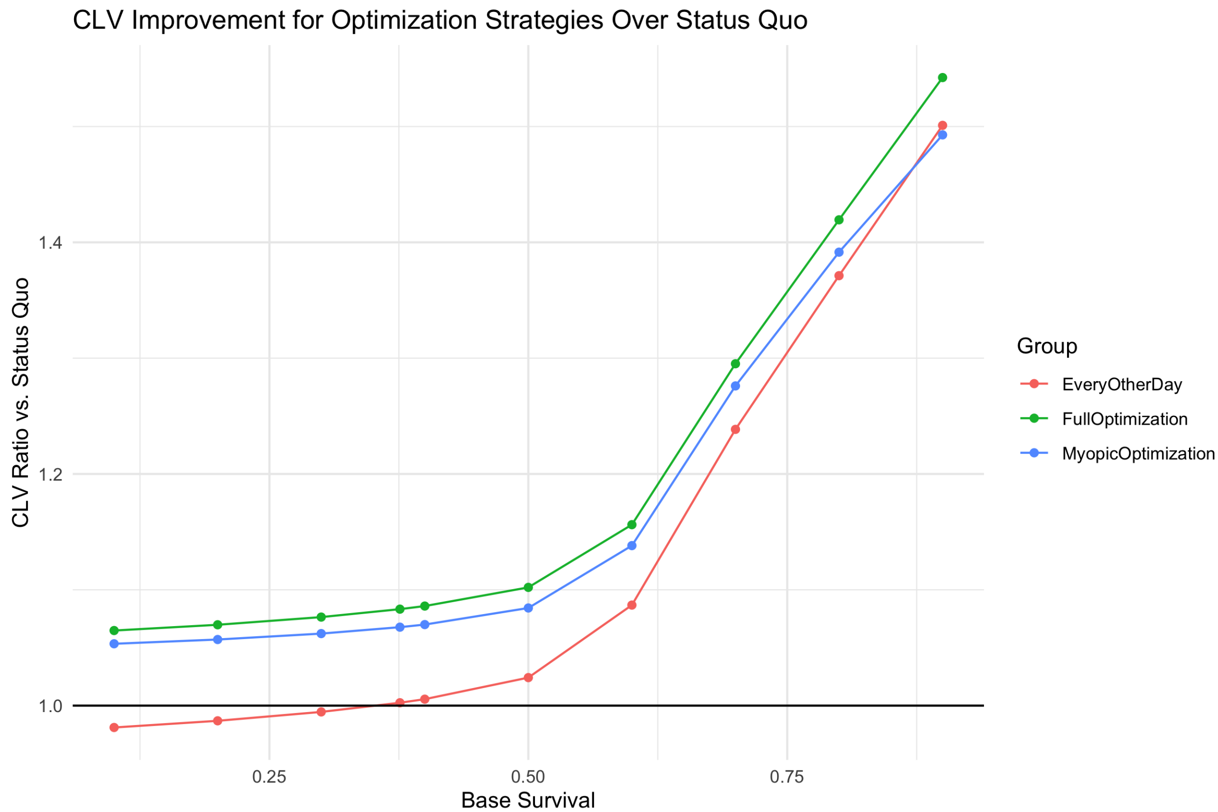

Figure 3 **Sending fewer emails is more valuable as base annual customer retention increases.**

Second, our findings highlight the value of customer segmentation and personalization when designing customer engagement policies. Many firms already personalize the content of emails or push notifications (e.g., product recommendations), but they often send those personalized messages on the same fixed schedule to everyone. Our research indicates substantial gains from personalizing the timing and frequency itself. Managers can use readily available information, such as customer tenure, prior purchasing activity, and email engagement (for example, open and click rates), to segment customers based on their sensitivity to email frequency. Even a simple segmentation (e.g., "high-engagement vs. low-engagement customers") could inform different email policies: highly engaged customers might paradoxically need fewer touches to keep them engaged and happy (since they actually read what you send and could be annoyed by too much), whereas disengaged customers can receive more frequent nudges without much risk (since many are ignoring them anyway). Our analysis specifically showed that new customers and customers who actively open emails benefited the most from reduced frequency in terms of retention, so those are prime candidates for a gentler touch. On the other hand, long-time loyalists and heavy spenders were relatively resilient to high frequency and generated more revenue from it, so the firm can maintain a high contact rate with that group. Implementing such differential treatment is feasible with modern email service providers, which allow segmentation and automated frequency control.

Third, the results encourage a mindset shift toward long-term metrics in marketing decision-making. We demonstrated how accounting for churn/retention fundamentally changes what is optimal. Campaign evaluations that look only at immediate customer spending might incorrectly conclude that more emails are always better, since they ignore churn as an outcome. Managers should incorporate metrics like unsubscribe rate, customer lifetime value, or at least retention-adjusted revenue when assessing policies designed to improve customer engagement on the platform. For instance, an email A/B test could measure not just conversion in the next week but also the downstream unsubscribe rate or customer value over the next quarter. In our field experiment, had we looked only at week-of revenue, the daily email strategy would appear superior; but by tracking unsubscribes and subsequent spend, we see the fuller picture. We recommend that firms routinely monitor customer unsubscribe (or spam complaint) rates alongside sales metrics for each campaign, and treat an uptick in unsubscribes as an indication that frequency or content strategy may need adjustment. Over time, building predictive models for customer churn risk as a function of contact history can help optimize customer engagement policies proactively.

Our findings align with broader industry moves toward more disciplined communication policies. For example, as it was mentioned above, LinkedIn publicly reported reducing the volume of emails it sends to members in response to user experience concerns, reflecting a growing recognition that aggressive contact strategies can backfire by accelerating disengagement (Collins 2015). A related lesson comes from aforementioned daily deal businesses such as Groupon, whose initial growth model relied on emailing a single deal each day to a subscriber list (Blumberg 2018). While this approach can look attractive through a purely short-run lens, our results quantify the long-run costs of pushing customers to opt out: opt-outs are not merely a lost channel, but a meaningful signal of annoyance that predicts lower subsequent purchasing and weaker retention. More broadly, the same logic plausibly extends beyond email to other direct marketing touchpoints, including mobile push notifications, SMS promotions, and retargeting ads. In each case, managers face a similar engagement annoyance tradeoff, where higher contact intensity can raise immediate response while also increasing the risk of opt-outs and long-run churn. The practical implication is that customer experience indicators, such as opt-outs and disengagement proxies, should be managed jointly with immediate conversion metrics. In operational terms, firms can combine causal estimates of short-run revenue effects with causal estimates of retention consequences to design frequency policies that maximize customer lifetime value rather than one-period customer spending.

A practical lesson from our counterfactual analysis is that the payoff to retention-aware contact policies is not universal: it depends on a firm's baseline churn and on whether customers tend to be one-time visitors or repeat purchasers. In high-retention businesses, customers have substantial remaining lifetime value, so preventing unsubscribe-driven churn is especially valuable. Even a modest reduction in opt-outs can preserve a long stream of future purchase opportunities, generating a disproportionate lift in customer lifetime value. In contrast, when baseline churn is high and many customers are effectively one-and-done, there is less future revenue at stake; consequently, the incremental benefit of protecting retention is smaller and short-run revenue considerations become relatively more important. This logic is reflected in our counterfactual results: the relative advantage of reduced frequency increases monotonically as baseline survival rises, and the CLV-optimal personalized policy delivers the largest gains throughout. For managers, the implication is to tailor "how hard you push" to the nature of the customer relationship. Firms with repeat purchasing and strong baseline loyalty should invest in retention-aware frequency caps and CLV-based personalization because the long-run upside dominates. By contrast, in high-churn settings, similar interventions will typically yield smaller profit gains.

6. Conclusion

The emergence of e-commerce and digital marketing has made it cheap and easy for firms to reach customers instantly and at scale. This creates a basic tension: sending messages more often can raise short-run spending, but it can also irritate customers, leading them to opt out and engage less over time. In this paper, we measure this engagement–annoyance tradeoff using a large randomized field experiment that varies coupon-email frequency, connect opt-out behavior to later purchasing, and turn the estimated response patterns into practical frequency policies. In this paper, we show that strategies chosen only to maximize immediate sales lift can backfire when annoyance builds and damages customer relationships in the long run.

Our findings offer several practical implications for managers who set and operate customer communication policies. First, email frequency should be treated as a core operating decision rather than a tactical adjustment. A higher sending rate can look attractive in the short run, yet become value-destroying once the downstream cost of additional opt-outs and weaker retention is taken into account. Second, opt-outs provide a direct operational signal of customer dissatisfaction with the communication process. Tracking opt-out rates by cohort and segment can flag when a policy is pushing customers beyond their tolerance and eroding the future effectiveness of the channel. Third, a single frequency rule is inefficient. The heterogeneous responses we document imply that some customer groups benefit substantially from fewer messages because opt-outs fall sharply, while other groups experience larger short-run spending losses when contact is reduced. Even simple segmentation based on tenure, prior purchase intensity, and historical engagement can support more stable, retention-aware communication rules. Finally, experimentation and evaluation should reflect long-run objectives. Because dissatisfaction and churn often appear with a delay, policy tests should be assessed not only on immediate revenue lift but also on opt-outs and subsequent purchasing, and these estimates should be used to set frequency rules that perform well over time.

Taken together, our results highlight a central managerial takeaway: promotional messages are not free simply because sending them is inexpensive. When annoyance is a binding constraint, the long-run value of customer attention must be explicitly incorporated into communication strategy, and retention-aware, personalized policies can outperform one-size-fits-all approaches.

References

Armony, Mor, Itai Gurvich. 2010. When promotions meet operations: Cross-selling and its effect on call-center performance. *Manufacturing & Service Operations Management* **12**(3) 470–488.

Awan, Aatif. 2015. Less is more: You're about to receive less email from linkedin. LinkedIn Blog. URL https://www.linkedin.com/blog/member/product/less-email-from-linkedin. Accessed: 2025-12-20.

Baardman, Lennart, Setareh Borjian Boroujeni, Tamar Cohen-Hillel, Kiran Panchamgam, Georgia Perakis. 2023. Detecting customer trends for optimal promotion targeting. *Manufacturing & Service Operations Management* **25**(2) 448–467.

Baardman, Lennart, Maxime C Cohen, Kiran Panchamgam, Georgia Perakis, Danny Segev. 2019. Scheduling promotion vehicles to boost profits. *Management Science* **65**(1) 50–70.

Baek, Jackie, Will Ma, Dmitry Mitrofanov. 2025. Personalized promotions in practice: Dynamic allocation and reference effects. *arXiv preprint arXiv:2512.23781* .

Besanko, David, Jean-Pierre Dubé, Sachin Gupta. 2003. Competitive price discrimination strategies in a vertical channel using aggregate retail data. *Management Science* **49**(9) 1121–1138.

Bitran, Gabriel R., Susana V. Mondschein. 1996. Mailing decisions in the catalog sales industry. *Management Science* **42**(9) 1364–1381.

Blumberg, Alex. 2018. The super-quick rise and even faster fall of groupon. Intelligencer (New York Magazine). URL https://nymag.com/intelligencer/2018/10/andrew-mason-on-groupon.html. Published October 10, 2018. Accessed December 20, 2025.

Chitla, Sandeep, Maxime C Cohen, Srikanth Jagabathula, Dmitry Mitrofanov. 2023. Customers' multihoming behavior in ride-hailing: Empirical evidence from uber and lyft. *Available at SSRN 4591826* .

Cohen, Maxime C, Jeremy J Kalas, Georgia Perakis. 2021. Promotion optimization for multiple items in supermarkets. *Management Science* **67**(4) 2340–2364.

Collins, Katie. 2015. Linkedin promises to quit bugging you so much. *WIRED* URL https://www.wired.com/story/linkedin-fewer-emails/. Accessed 2026-01-30.

Cook, John. 2014. Email delivery snafus cause headaches for zulily in q3, company cuts marketing spending as it works through issues. GeekWire. URL https://www.geekwire.com/2014/email-delivery-snafus-cause-headaches-zulily-q3-company-cuts-marketing-spending-works-issues/. Accessed: 2025-12-20.

Fader, Peter S., Bruce G. S. Hardie, Ka Lok Lee. 2005. Rfm and clv: Using iso-value curves for customer base analysis. *Marketing Science* **24**(2) 275–284.

Fong, Nathan M. 2017. How targeting affects customer search: A field experiment. *Management Science* **63**(7) 2353–2364.

Jagabathula, Srikanth, Dmitry Mitrofanov, Gustavo Vulcano. 2022. Personalized retail promotions through a directed acyclic graph–based representation of customer preferences. *Operations Research* **70**(2) 641–665.

Kadiyala, Bharadwaj, Özalp Özer, A Serdar Şimşek. 2024. Dual value of delayed incentives: An empirical investigation of gift card promotions. *Manufacturing & Service Operations Management* **26**(5) 1806–1825.

Lin, Yunduan, Mengxin Wang, Heng Zhang, Renyu Zhang, Zuo-Jun Max Shen. 2024. Content promotion for online content platforms with the diffusion effect. *Manufacturing & Service Operations Management* **26**(3) 1062–1081.

Liu, Xiao. 2023. Dynamic coupon targeting using batch deep reinforcement learning: An application to livestream shopping. *Marketing Science* **42**(4) 637–658.

Murthi, B. P. S., Sumit Sarkar. 2003. The role of the management sciences in research on personalization. *Management Science* **49**(10) 1344–1362.

Reimers, Imke, Claire Xie. 2019. Do coupons expand or cannibalize revenue? evidence from an e-market. *Management Science* **65**(1) 286–300.

Reinartz, Werner J., V. Kumar. 2003. The impact of customer relationship characteristics on profitable lifetime duration. *Journal of Marketing* **67**(1) 77–99.

Sahni, Navdeep S., Dan Zou, Pradeep K. Chintagunta. 2017. Do targeted discount offers serve as advertising? evidence from 70 field experiments. *Management Science* **63**(8) 2688–2705.

Schonfeld, Erick. 2011. Closing the redemption loop in local commerce. TechCrunch. URL https://techcrunch.com/2011/07/24/redemption-loop-local-commerce/. Accessed: 2025-12-20.

Shaffer, Greg, Z. John Zhang. 2002. Competitive one-to-one promotions. *Management Science* **48**(9) 1143–1160.

Todri, Vilma, Anindya Ghose, Param Vir Singh. 2020. Trade-offs in online advertising: Advertising effectiveness and annoyance dynamics across the purchase funnel. *Information Systems Research* **31**(1) 102–125. doi:10.1287/isre.2019.0877. URL http://dx.doi.org/10.1287/isre.2019.0877.

Venkatesan, Rajkumar, V. Kumar. 2004. A customer lifetime value framework for customer selection and resource allocation strategy. *Journal of Marketing* **68**(4) 106–125.

Wang, Yiwei, Lauren Xiaoyuan Lu, Pengcheng Shi. 2022. Does customer email engagement improve profitability? evidence from a field experiment in subscription service retailing. *Manufacturing & Service Operations Management* **24**(5) 2703–2721.

Weber, Harrison. 2013. Fab stops sending you emails you don't read, even when you don't ask them to. The Next Web. URL https://thenextweb.com/news/classy-fab-stops-sending-you-emails-you-dont-read-even-when-you-dont-ask-them-to. Accessed: 2025-12-20.

Appendix A: Robustness Checks: Log Transform

We show that the main results are robust to log transforming key dependent variables. Table 14 shows that after log transforming revenue, item quantity, and price per item, the results in Table 2 remain consistent: sending fewer emails decreases short-term revenue and item quantity but does not affect price per item purchased.

Table 14 Average Treatment Effects for Reducing Email Frequency, Log Transform

	Dependent variable:		
	log(1+Revenue)	log(1+Item Quantity)	log(Price Per Item)
	(1)	(2)	(3)
Treated	-0.001^{***}	-0.0001^{***}	-0.010
	(0.0002)	(0.00005)	(0.019)
Observations	4,222,071	4,222,071	9,635
Fixed Effects	Date	Date	Date
Note:			*p<0.1; **p<0.05; ***p<0.01

Additionally, from the historical data, Table 15 shows that after log transforming revenue, the results in Table 7 remain consistent.

Table 15 Unsubscribing Significantly Decreases Revenue, Log Transform

	Dependent variable:	
	log(1+Month Revenue)	
	(1)	(2)
Previous Month Revenue ($)	0.0002^{***}	0.0002^{***}
	(0.00001)	(0.00001)
Previous Month Unsubscribed	-0.174^{***}	-0.188^{***}
	(0.008)	(0.011)
Observations	1,095,832	1,095,832
Fixed Effects	Month	Month
IV	None	Previous Month Email Frequency
Note:		*p<0.1; **p<0.05; ***p<0.01

Appendix B: Robustness Checks: Model Specifications

We show that the results in Table 2 are robust to alternative model specifications beyond the linear model. Specifically, we use logistic regression for the probabilities of unsubscribing and of purchasing items and Poisson regression for item quantity. The results remain consistent in showing that sending fewer emails generates a tradeoff: lower email frequency improves retention by decreasing the probability of unsubscribing but also reduces short-term revenue by decreasing the probability of customers making a purchase and the average quantity of items purchased. Table 16 shows the results.

Table 16 Average Treatment Effects for Reducing Email Frequency, Alternative Model Specifications

	Dependent variable:		
	P(Unsubscribe)	Item Quantity	P(Purchase)
	(1)	(2)	(3)
Treated	-0.896^{***}	-0.050^{**}	-0.105^{***}
	(0.069)	(0.020)	(0.015)
Constant	-7.974^{***}	-6.057^{***}	-5.463^{***}
	(0.037)	(0.014)	(0.010)
Observations	4,222,071	4,222,071	4,222,071
Model	Logistic	Poisson	Logistic
Note:			*p<0.1; **p<0.05; ***p<0.01

Additionally, from the historical data, Table 17 shows that the results for customer retention in Table 8 are robust to logistic regression.

Table 17 Effect of Unsubscribing on Customer Retention, Alternative Model Specification

	Dependent variable:	
	P(Survival)	
	After Month 1	After Months 1-6
	(1)	(2)
Unsubscribed	-1.157^{***}	-1.002^{***}
	(0.275)	(0.063)
Tenure	0.0002^{***}	0.0001^{***}
	(0.00001)	(0.00000)
Unsubscribed×Tenure	-0.00004	0.00001
	(0.0001)	(0.00002)
Constant	-1.474^{***}	-2.363^{***}
	(0.023)	(0.014)
Observations	23,465	118,941
Model	Logistic	Logistic
Note:		*p<0.1; **p<0.05; ***p<0.01

Appendix C: Robustness Checks: Random Assignment and Difference-in-Differences

For the field experiment, we confirm that users are randomly assigned to the treatment group in Table 18 by comparing user features in the week before the experiment. In the user features, we omit unsubscribing as users who unsubscribe before the experiment no longer receive emails and are therefore excluded from the experiment. Across the variables, differences are statistically insignificant, confirming that users in the treatment and control groups have similar characteristics.

Table 18 Confirming Random Assignment: Similar Average Values for Control and Treated Users in Week Before Experiment

	Revenue (\$/Day)	Item Quantity	P(Purchase)	Price Per Item
Control	0.3555	0.0030	0.0017	132.09
Treated	0.3466	0.0029	0.0018	131.01
Rel. Size	-2.5%	-1.7%	+3.0%	-0.8%
t-test p-value	0.597	0.703	0.208	0.809
Note:				*p<0.1; **p<0.05; ***p<0.01

In addition to confirming random assignment, we show that the main results in Table 2 are robust to difference-in-differences (DiD) in Table 19. In the DiD specification, we include customer fixed effects, yielding the specification

$$\text{Outcome}_{it} = \text{Treated}_i + \text{Post}_t + \text{Treated}_i \times \text{Post}_t \mid \text{Customer}_i + \text{Date}_t \tag{8}$$

All results are directionally consistent, and significance is also consistent except that the negative effect on revenue becomes insignificant. However, a smaller or no revenue loss would keep our overall results consistent, as sending fewer emails would be more valuable under the same benefit to customer retention. Additionally, Table 20 shows the DiD specification after log transforming numeric variables, yielding consistent results with the same significance as Table 2.

Table 19 Average Treatment Effects for Reducing Email Frequency, Difference-in-Differences

	Dependent variable:				
	P(Unsubscribe)	Revenue (\$/Day)	Item Quantity	P(Purchase)	Price Per Item
	(1)	(2)	(3)	(4)	(5)
Treated×Post	-0.0002^{***}	-0.028	-0.0004^{***}	-0.0002^{***}	38.619
	(0.00002)	(0.024)	(0.0001)	(0.0001)	(34.295)
Observations	4,222,071	4,222,071	4,222,071	4,222,071	9,635
Fixed Effects			Customer + Date		
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 20 Average Treatment Effects for Reducing Email Frequency, Difference-in-Differences with Log Transform

	Dependent variable:		
	log(1+Revenue)	log(1+Item Quantity)	log(Price Per Item)
	(1)	(2)	(3)
Treated	-0.001^{***}	-0.0002^{***}	0.046
	(0.0003)	(0.0001)	(0.100)
Observations	8,444,142	8,444,142	16,992
Fixed Effects	Date	Date	Date
Note:			*p<0.1; **p<0.05; ***p<0.01

Appendix D: Robustness Checks: Binary Transform for HTEs

We show that the HTE results in Section 4 are robust to binary transforming the HTE variables on a median split. Tables 21 to 24 show the corresponding results remain consistent for Tables 3 to 6. In particular, the interaction term coefficients show consistent patterns: customer tenure and email open rates continue to significantly predict unsubscription, while recent purchase behavior remains the primary driver of revenue.

Table 21 HTEs: Customer Tenure, Median Split

	Dependent variable:	
	P(Unsubscribe)	Revenue ($/Day)
	(1)	(2)
Tenure>2577 Days	-0.0001^{**}	0.025
	(0.00002)	(0.025)
Treated	-0.0002^{***}	-0.050^{**}
	(0.00002)	(0.025)
Tenure>2577 Days×Treated	0.00005	0.027
	(0.00003)	(0.036)
Observations	4,222,071	4,222,071
Note:		*p<0.1; **p<0.05; ***p<0.01

Additionally, from the historical data, Table 25 shows consistent results with Table 9, with additional significant interaction effects for customer tenure and email opens and a consistent positive interaction effect on item quantity.

Appendix E: Robustness Checks: Changes to Instrumental Variables

We show that the instrumental variables results in Table 7 is robust to two alternative specifications for the instrumental variable of past month email frequency. First, we use logit in the first stage in Table 26. Second,

Table 22 HTEs: Item Quantity in Past 7 Days, Median Split

	Dependent variable:	
	P(Unsubscribe)	Revenue ($/Day)
	(1)	(2)
Item Quantity>0	0.0003***	1.807***
	(0.0001)	(0.117)
Treated	−0.0002***	−0.022
	(0.00002)	(0.018)
Item Quantity>0×Treated	−0.00002	−1.250***
	(0.0001)	(0.166)
Observations	4,222,071	4,222,071
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 23 HTEs: Email Opens in Past 7 Days, Median Split

	Dependent variable:	
	P(Unsubscribe)	Revenue ($/Day)
	(1)	(2)
Email Opens>0	0.0004***	0.025
	(0.00002)	(0.025)
Treated	−0.0001***	−0.031
	(0.00002)	(0.026)
Email Opens>0×Treated	−0.0003***	−0.012
	(0.00003)	(0.036)
Observations	4,222,071	4,222,071
Note:		*p<0.1; **p<0.05; ***p<0.01

we use a leave-one-out version in Table 27, in which the instrument is the average email frequency of other customers in the previous month. The results are consistent in both specifications.

Additionally, we confirm that the instrument is relevant in Table 28, with a first stage F-statistic of 1,498,678.

Table 24 HTEs: All Variables Combined, Median Split

	Dependent variable:	
	P(Unsubscribe)	Revenue (\$/Day)
	(1)	(2)
Tenure>2577 Days	-0.0001^{***}	0.024
	(0.00002)	(0.025)
Item Quantity>0	0.0003^{***}	1.806^{***}
	(0.0001)	(0.117)
Email Opens>0	0.0004^{***}	0.022
	(0.00002)	(0.025)
Treated	-0.0001^{***}	−0.031
	(0.00003)	(0.032)
Tenure>2577 Days×Treated	0.0001^{*}	0.027
	(0.00003)	(0.036)
Item Quantity>0×Treated	−0.00001	-1.250^{***}
	(0.0001)	(0.166)
Email Opens>0×Treated	-0.0003^{***}	−0.010
	(0.00003)	(0.036)
Observations	4,222,071	4,222,071
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 25 HTEs for Effect of Unsubscribing on Revenue, Median Split

	Dependent variable:		
	Month Revenue ($)		
	(1)	(2)	(3)
Previous Month Revenue	0.087^{***}	0.086^{***}	0.112^{***}
	(0.001)	(0.001)	(0.001)
HTE	6.103^{***}	2.647^{***}	-27.759^{***}
	(0.246)	(0.254)	(0.399)
Previous Month Unsubscribed	-5.281^{***}	-25.007^{***}	-11.593^{***}
	(1.371)	(2.935)	(1.071)
Previous Month Unsubscribed×HTE	-5.364^{***}	17.209^{***}	9.527^{***}
	(1.958)	(3.116)	(2.516)
Observations	1,095,832	1,095,832	1,095,832
HTE	Tenure	Email Opens	Item Quantity
Fixed Effects	Month	Month	Month
IV	Yes	Yes	Yes
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 26 IV, Logit First Stage

	Dependent variable:
	Month Revenue ($)
Previous Month Revenue	0.086^{***}
	(0.001)
Previous Month Unsubscribed	-0.074^{***}
	(0.009)
Observations	1,095,832
Instrument	Previous Month Email Frequency
First Stage Model	Logit
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 27 IV, Leave One Out

	Dependent variable:
	Month Revenue ($)
Previous Month Revenue	0.086^{***}
	(0.001)
Previous Month Unsubscribed	-8.081^{***}
	(0.980)
Observations	1,095,832
Instrument	Other Customers' Previous Month Email Frequency
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 28 First Stage for Linear IV Model

	Dependent variable:
	Previous Month Unsubscribed
Previous Month Email Frequency	0.024^{***}
	(0.00002)
Constant	-0.026^{***}
	(0.0001)
Observations	1,095,832
F-Statistic	1,498,678
Note:	*p<0.1; **p<0.05; ***p<0.01