

<b>Abstract</b>	<b>4</b>
<b>1. Introduction</b>	<b>5</b>
1.1 The return crisis in fashion e-commerce	5
1.2 Traditional return handling is broken	6
1.3 The gap in the literature	7
1.4 Research question and contribution	9
<b>2. Literature Review</b>	<b>11</b>
2.1 Digital nudging and choice architecture	11
2.1.1 Nudging as a concept	11
2.1.2 From physical to digital nudging	12
2.1.3 Digital nudging in e-commerce returns - the frontier	13
2.2 Consumer behaviour in the return decision	14
2.2.1 Cognitive dissonance and buyer's remorse	14
2.2.2 Autonomy and self-determination in designed environments	15
2.2.3 Moral licensing and the downstream effect of virtuous choice	16
2.3 Sustainability, circular economy, and the rebound effect	17
2.3.1 Rebound in the circular economy	17
2.3.2 Consumer-side rebound	18
2.3.3 Fashion returns as an environmental problem	19
<b>3. Case Presentation</b>	<b>21</b>
3.1 Returnflows	21
3.2 The return flow	21
3.3 The value proposition	23
3.4 The empirical case	24
<b>4. Methodology</b>	<b>26</b>
4.1 Philosophy of science	26
4.2 Research design	26
4.3 Data collection and sources	27
4.4 Analysis, quality criteria and limitations	28
<b>5. Analysis and Discussion</b>	<b>30</b>
5.1 The choice architecture: how Returnflows nudges the return decision	30

5.1.1 Exchange as the default: visual hierarchy and salience	30
5.1.2 The refund fee as loss aversion nudge	31
5.1.3 The gift card bonus as gain framing	32
5.1.4 Friction design: instant exchange and the cost of refund	33
5.1.5 Synthesis: the architecture as a system	34
5.2 Behavioural outcomes: what happens to consumer decisions and merchant value	34
5.2.1 From a single endpoint to a three-way split	34
5.2.2 The strategic value of the redistribution	35
5.2.3 The consumer experience: autonomy or manufactured choice?	36
5.3 The sustainability paradox: rebound risk at the point of return	38
5.3.1 The case for net environmental benefit	38
5.3.2 The rebound counterargument	39
5.3.3 Synthesis: where does the evidence land?	41
<b>6. Conclusion</b>	<b>43</b>
6.1 Answering the research question	43
6.2 Limitations	44
6.3 Recommendations and perspectives	44
Literature	45
A strategic briefing on the Returnflows choice architecture	49
The problem and the opportunity	49
What the architecture does, and what it has produced	50
Strategic recommendations	51
Closing note	52

**COPENHAGEN BUSINESS SCHOOL**

BSc in Business Administration and Communication - HA(kom.)

Bachelor Thesis · Spring 2026

## Abstract

Returns are the structural weak point of fashion e-commerce. In the United States alone, 2025 retail returns totalled 849.9 billion dollars, with online return rates reaching 19.3%, and fashion sits at the top of that distribution. Yet most merchants still process every return through a single default path on platforms such as Shopify: the customer requests a refund, receives a label, sends the garment back, and the transaction ends as a loss. Each return processed this way represents lost revenue plus an outbound reverse-logistics cost, and these losses now compound as Extended Producer Responsibility legislation in Europe begins to add the end-of-life cost of discarded garments to the merchant's balance sheet.

This thesis examines an emerging alternative. Returnflows is a SaaS platform that replaces the default Shopify refund flow with a three-option choice architecture: exchange, store credit with a bonus, or refund. The analytical framework combines three streams. Digital nudging and choice-architecture theory (Thaler & Sunstein, 2008; Weinmann et al., 2016) explains how the interface guides the return decision. Cognitive dissonance (Festinger, 1957) and self-determination theory (Ryan & Deci, 2000) interpret why consumers respond to the design, and whether their choice remains genuinely autonomous. Circular economy rebound theory (Zink & Geyer, 2017; Ciecelska et al., 2024) tests whether redirecting returns away from refund produces environmental gains that survive downstream consumer behaviour.

The research design is an abductive case study grounded in pragmatism. Empirical material comprises quantitative platform data from a Danish fashion merchant covering December 2023 to March 2026, a semi-structured interview with the Returnflows founder, and secondary industry and regulatory sources.

Three findings anchor the analysis. First, the choice architecture shifts the outcome distribution from a near-universal refund baseline to 62.27% refund, 30.08% exchange, and 7.65% store credit by return value. Of 7.32 million DKK in total return value, 3.45 million DKK is retained, a retained ratio of 47.16%. Second, 77.5% of returns are size or fit related, which reframes the typical return as a variant-correction problem rather than a rejection of the purchase, and maps onto dissonance-reduction logic that predicts exchange as the psychologically cheapest resolution. Third, the environmental case is strongest where the

rebound literature is weakest. Exchange replaces two reverse-logistics trips with one, draws from already-committed demand rather than stimulating new consumption, and feeds a data loop that reduces size driven returns upstream.

The conclusion is qualified. The choice architecture works as designed, the commercial case against the Shopify baseline is straightforward, and the sustainability case is net positive when limited to the exchange pathway. Rebound risk is real but structurally contained in the store credit pathway. Two factors bound the risk empirically: the 27-day average gift card usage window, suggesting purposeful rather than impulsive deployment, and the bounded demand structure of a post-purchase context, where the consumer's intent to own the product has already been settled.

The contribution is threefold, mapped to the HA(kom.) disciplines. To Communication, the thesis extends digital nudging theory into a post-purchase context that existing research has largely ignored. To Business Administration, it quantifies the strategic value of return redesign against the natural platform baseline and reframes returns as a revenue-retention mechanism. To Social Diagnostics, it examines consumer behaviour and environmental consequence in the same analytical frame, asking whether a design that resolves dissonance efficiently can do so without producing rebound downstream.

## **1. Introduction**

### **1.1 The return crisis in fashion e-commerce**

Few aspects of contemporary retail have scaled as quickly, and as uncomfortably, as the online return. In 2024, U.S. retailers absorbed 849.9 billion USD in returned merchandise, with roughly one in five online purchases sent back (NRF & Happy Returns, 2025). Fashion sits at the sharp end. Marriott et al. (2025) estimate that returns cost UK fashion retailers 7 billion GBP in 2022 alone, with category-level return rates running at 34 to 36% during peak periods. Stocker et al. (2021) report return rates of up to 50% in German online fashion, while Nestler et al. (2021) document categories touching 75%. The industry has built a parallel supply chain moving backwards.

These factors are systemic. Clothing is what Diggins et al. (2016) term a high-risk purchase: buyers cannot feel fabric weight, test fit, or see true colour until the parcel arrives, and the gap between expectation and reality drives the return. Stocker, Baier and Brand (2021) confirm 87.1% of returns are fit-related, with 41.6% of consumers admitting to bracketing. The infrastructure built around this scale is now itself a source of cost: handling, transport, restocking, write-downs and disposal, with reverse logistics costing up to nine times the forward equivalent.

This was already an expensive problem when margins were healthy. It has become existential as they compress. The State of Fashion 2026 report finds that 75% of fashion executives plan to raise prices in 2026, while 80% of consumers say they will refuse to buy at full price (Business of Fashion & McKinsey, 2026). Inventory days outstanding have risen 14% to 168 days, with reverse logistics costing up to nine times the forward equivalent (Diggins et al., 2016). Calcabrini et al. (2020) suggest retailer profits could rise by approximately 50% if returns were eliminated, illustrating how much margin the category absorbs in the return process. Layered on top is the European Extended Producer Responsibility regime: France imposes per-unit fines of up to 7,500 EUR, and the EU Ecodesign for Sustainable Products Regulation extends comparable obligations across the bloc.

This was already an expensive problem when margins were healthy. It has become existential as they compress. The State of Fashion 2026 report finds that 75% of fashion executives plan to raise prices in 2026, while 80% of consumers say they will refuse to buy at full price (Business of Fashion & McKinsey, 2026). Inventory days outstanding have risen 14% to 168 days, with reverse logistics costing up to nine times the forward equivalent (Diggins et al., 2016). Calcabrini et al. (2020) suggest retailer profits could rise by approximately 50% if returns were eliminated, illustrating how much margin the category absorbs in the return process.

## 1.2 Traditional return handling is broken

Against this backdrop, the way many fashion webshops actually process returns is surprisingly under-optimised. Shopify includes native support for exchanges and store credit, and its self-serve portal typically asks the customer to select a return reason. The infrastructure is there. But enabling it, designing the flow so exchange and store credit are surfaced as real alternatives rather than buried options, and acting on the reason-code data, are separate decisions many merchants have not made. Without deliberate design at the resolution screen, the default path most customers take remains the refund.

This default reflects a particular commercial philosophy. Returns were historically treated as a cost of doing business online, to be minimised where possible and absorbed where not. The operational focus has been on reducing return volume, through better size guides, more accurate product photography, and clearer policies, rather than on reshaping the return event itself. Yet the customer in a return context is not necessarily exiting the relationship. A customer returning a jumper because it was too small still wants a jumper. The current default treats every return as the same problem, when it is not.

That framing misses something important. A consumer who returns a jumper that was too small does not necessarily want a refund. They want a jumper that fits. Offering only a refund resolves the transaction but not the underlying need, and in the process hands the subsequent purchase to a competitor. Traditional return handling, in this sense, is not merely inefficient. It is strategically inert. It processes the customer's immediate demand without asking whether that demand reflects their actual goal.

A new generation of software has begun to challenge this default. Platforms such as Returnflows, Loop Returns, and Narvar redesign the return journey as a structured decision environment in which exchange, store credit, and refund are presented as parallel options, each framed differently and each carrying different consequences for the consumer. The return becomes a choice architecture rather than a fixed pipeline. Whether that architecture guides consumers toward better outcomes, or merely toward outcomes more profitable for the retailer, is the question this thesis takes up.

### 1.3 The gap in the literature

Nudging theory provides a natural starting point. Thaler and Sunstein (2008) established that small changes in how choices are presented produce large changes in what people choose, without restricting freedom or altering economic incentives. Weinmann et al. (2016) translated the framework to digital environments. The empirical literature that followed has been substantial, but its centre of gravity sits in the pre-purchase phase. Product pages, checkout flows, and search interfaces have been studied in detail. The post-purchase journey, and the return decision specifically, remains comparatively uncharted.

The few studies that address returns directly confirm the gap rather than close it. Von Zahn et al. (2024) demonstrate that machine-learning-targeted green nudges during checkout reduce return shipments by 2.6% without harming sales. Rauh et al. (2024) show that a single green image during purchase reduces return motivation by nearly a full point on a seven-point scale. Both are pre-purchase interventions on the prevention side. Neither examines what happens at the moment a consumer has already decided to return an item and faces a structured choice between refund, exchange, and store credit. Rauh et al. (2024) note that “the potential use of digital nudging to decrease return motivation in particular has received limited research attention so far” (p. 2).

The consumer-behaviour literature on returns is richer but disconnected from design. Festinger’s (1957) theory of cognitive dissonance explains why a mismatched purchase generates psychological discomfort the consumer is motivated to resolve, but it was developed long before anyone had to choose between a refund button and an exchange button on a smartphone. Ryan and Deci’s (2000) self-determination theory offers a vocabulary for thinking about autonomy in designed choice environments, yet its application to consumer e-commerce is thin. Blanken et al. (2015) show that virtuous acts can license subsequent indulgence, a mechanism directly relevant to whether choosing exchange makes the next purchase more careless, but moral licensing has not been tested in return flows.

A parallel gap sits in the sustainability literature. Zink and Geyer (2017) warn that profitable circular-economy innovations tend to produce rebound effects, and Yerushalmi and Saha (2025) quantify a 155% rebound backfire in the textiles and clothing sector. These are production-side analyses. Ciechelska et al. (2024) and Makov and Meshulam (2025) bring

the question to the consumer, but in secondhand markets rather than returns. No study, to my knowledge, asks whether redirecting a consumer from refund to exchange produces a net environmental gain, a net loss, or something in between once downstream behaviour is accounted for.

The result is a triangular gap. Nudging is studied in pre-purchase contexts but not in return flows. Return psychology is studied as individual behaviour but not as a response to designed choice architecture. Rebound is studied at the production and secondhand-market level but not at the level of the return decision itself. Returnflows, as a live system redesigning the return journey for Shopify fashion merchants, sits precisely at the intersection of all three.

#### **1.4 Research question and contribution**

This thesis investigates how that intersection works in practice. The central research question is:

***How does Returnflows use digital design and nudging to guide consumer behaviour during apparel and fashion returns, and what strategic value does this offer webshops compared to traditional return handling?***

Four sub-questions structure the analysis. The first asks how Returnflows uses digital design and nudging in the online return journey, drawing on the communication and choice-architecture literature to describe the mechanism. The second asks how this design guides or changes consumer behaviour during returns, combining cognitive-dissonance and self-determination theory to interpret the observed shift from refund-dominated outcomes toward exchange and store credit. The third asks what strategic value the system creates compared to traditional return handling, quantifying revenue retention against the near-universal refund baseline that Shopify's customer-facing default produces. The fourth asks whether guiding consumers away from refunds carries an environmental cost through the rebound effect.

The contribution is threefold, mapped to the three disciplines that define the HA(kom.) programme. To Communication, the thesis extends digital nudging theory into a post-purchase context that existing research has largely ignored, and documents how a

coherent choice architecture is built from individually familiar elements: defaults, framing, friction, and incentives. To Business Administration, it quantifies the strategic value of redesigning returns against the natural Shopify baseline, reframing the return from a cost centre into a revenue-retention mechanism. To Social Diagnostics, it examines the behavioural and environmental consequences in the same frame: whether a design that resolves dissonance efficiently and respects perceived autonomy can nonetheless produce rebound downstream. The 77.5% size-driven return rate matters here, because it shows that most returns are not about rejecting the product but about getting the wrong size.

The thesis proceeds in six sections. Section 2 develops the theoretical framework across three streams: digital nudging, consumer behaviour during returns, and rebound in circular systems. Section 3 presents Returnflows as a case. Section 4 sets out the methodology. Section 5 integrates the three disciplinary frameworks through a structured analysis of the choice architecture, the behavioural outcomes it produces, and the sustainability paradox it raises. Section 6 concludes with an answer to the research question, the study's limitations, and recommendations for Returnflows, the broader e-commerce sector, and future research.

## 2. Literature Review

### 2.1 Digital nudging and choice architecture

#### 2.1.1 *Nudging as a concept*

The modern vocabulary of behavioural design begins with Thaler and Sunstein (2008). Their starting point is unfashionable in economics but obvious to anyone who has designed a form: people do not always act as the rational agents classical theory assumes. Attention is finite, willpower inconsistent, and decisions are made under cognitive shortcuts that produce predictable mistakes. Thaler and Sunstein treat this as a responsibility of the person designing the choice. Their term for that person is the choice architect, and their term for the small, non-coercive design decisions that shape outcomes is the nudge.

Thaler and Sunstein (2008) define a nudge as “any aspect of the choice architecture that alters people’s behavior in a predictable way without forbidding any options or significantly changing their economic incentives” (p. 6). The definition is narrower than it looks. A nudge preserves freedom of choice: all options remain available, and the cost of any of them stays essentially unchanged. Placing fruit at eye level nudges. Banning dessert does not. The framework they call libertarian paternalism accepts that someone has to design the environment in which a choice is made, that there is no neutral default, and that the architect may as well design toward outcomes that serve the chooser’s own interests.

The mechanisms are by now well catalogued. Defaults exploit the fact that most people accept whatever option is pre-selected, a tendency Thaler and Sunstein (2008) trace to inattention and status quo bias. Loss aversion, the finding that people weigh losses roughly twice as heavily as equivalent gains, turns small fees into powerful deterrents. Framing shifts behaviour by changing the linguistic packaging of an identical outcome. Friction, the effort tax attached to a given action, makes desirable behaviours easier and undesirable ones harder without prohibiting anything. Empirically, cafeteria rearrangements shift consumption by up to 25%, opt-out organ donor defaults nearly double consent, and the fly etched into the Schiphol urinal reduced spillage by around 80%. None of these restricted behaviour. All of them changed it.

### ***2.1.2 From physical to digital nudging***

Nudging was developed in a world of cafeterias and paper forms. Digital environments operate under different constraints and offer different levers. Weinmann et al. (2016) provide the conceptual bridge, defining digital nudging as “the use of user-interface design elements to guide people’s behavior in digital choice environments” (p. 433). The reformulation matters. In a digital context, every pixel is a design decision. Button size, colour, vertical order, click depth, the copy on a tooltip: each is a choice the architect makes, and each influences the outcome whether anyone intended it to.

The central claim Weinmann et al. (2016) draw from this is that digital neutrality is impossible. Since there is no neutral way to present choices in a user interface, every design decision influences users’ behaviour, frequently regardless of what the designer intended. A checkout flow that asks for a newsletter opt-in above the purchase button nudges. So does one that asks below. So does one that does not ask at all. The question is not whether to nudge but whether to nudge deliberately.

Digital environments also amplify the nudge toolkit. Weinmann et al. (2016) catalogue a set of principles adapted from the offline literature: defaults, incentives, feedback, error tolerance, structuring complex choices, and helping users understand mappings between options and outcomes. What separates the digital application from the offline one is the speed and detail of iteration. An airline can A/B test two checkout flows across hundreds of thousands of users in a week. A cafeteria cannot. Personalisation, machine-learning targeting, and real-time feedback loops make digital nudging a moving architecture rather than a fixed one.

Weinmann et al. (2016) are explicit about the ethical stakes. Nudging was conceived as a way to help people make better choices for themselves, and in commercial digital environments that condition cannot be taken for granted. They cite European low-cost carriers that use default-based nudges to push customers toward non-essential purchases, noting that such interventions may produce short-term gains but “long-term repercussions in terms of loss of goodwill, negative publicity, or even legal action” (p. 434). The designer’s responsibility is not merely to nudge effectively but to nudge in a direction the user would,

on reflection, endorse. This ethical frame will matter when the thesis turns to whether Returnflows' architecture supports or undermines consumer autonomy.

### ***2.1.3 Digital nudging in e-commerce returns***

Most of the empirical work on digital nudging in commerce has concentrated on the pre-purchase phase. Product pages, search results, checkout flows, and onboarding screens have been tested with conversion rate optimization (CRO) extensively. The return journey, by contrast, remains comparatively uncharted. Two recent studies have begun to close that gap, and the way they define its edges is informative.

Von Zahn et al. (2024) conducted a field experiment with a European fashion retailer involving 117,304 customers. The intervention was a shopping-cart prompt drawing attention to the environmental cost of bracketing. The result was a 2.6% reduction in return shipments with no adverse effect on sales, translating to roughly 340,000 USD in annual savings for the partner firm. Targeting the nudge via causal machine-learning roughly doubled effectiveness. The intervention operates pre-purchase, and its effects were heterogeneous: a backfiring pattern in approximately 40% of customers undercut aggregate gains.

Rauh et al. (2024) tested whether a single green image alongside a product reduces return motivation in a survey of 988 respondents. It does, by approximately one full point on a seven-point scale, but it also lowered purchase motivation through measurable reductions in autonomy (from 5.45 to 4.97) and competence (from 4.98 to 4.53). The authors note that a nudge aimed at one behaviour can bleed into adjacent ones, with the green-image cue reducing both return motivation and the autonomy and competence consumers feel during purchase.

Two findings from this literature are central. First, digital nudging in returns has been studied almost exclusively as a prevention problem, with interventions placed before the purchase. Second, pre-purchase nudges carry a structural risk of spilling over into purchase motivation, because they ask the consumer to doubt the wisdom of buying at the moment they are deciding whether to buy. Neither applies cleanly to a post-purchase intervention that

operates after the consumer has already decided the item is going back. At that point, the purchase motivation has already done its work. The question is no longer whether to buy but what form of resolution to accept. This is the design space Returnflows occupies, and the space the existing literature has not yet addressed.

*Nudging theory explains how digital choice architectures influence behaviour at the level of mechanism. It does not, on its own, explain why consumers respond to these architectures during returns. For that, the thesis turns to the psychology of the return decision itself: the dissonance that triggers the return, the autonomy the consumer brings to it, and the licensing effects that may follow.*

## **2.2 Consumer behaviour in the return decision**

### **2.2.1 Cognitive dissonance and buyer's remorse**

Nudging describes how choice architectures shape behaviour. It does not explain why anyone is returning a garment in the first place. For that, the thesis turns to Festinger (1957), whose theory of cognitive dissonance is the foundational account of why expectation-violating experiences are not merely disappointing but psychologically motivating. Festinger's central hypothesis is that "the existence of dissonance, being psychologically uncomfortable, will motivate the person to try to reduce the dissonance and achieve consonance" (p. 3). The claim applies directly to the moment a consumer opens a package and finds that the jumper is too small.

Festinger (1957) Festinger (1957) treats dissonance as a motivational state comparable to hunger or frustration, a drive that presses the person toward any behaviour that restores consonance. For an online purchase, the relevant cognitions are straightforward: the consumer chose this garment, ordered this size, and expected it to fit. When it does not, the tension demands resolution.

Festinger (1957) identifies three reduction strategies, each of which maps onto a branch of the Returnflows return flow. The consumer can change the behavioural element, which in this context means reversing the purchase: returning the garment for a refund. The consumer can change the environmental element, the variant of the product that produced the dissonance: exchanging the wrong size for the right one. Or the consumer can add new

consonant elements that reduce dissonance without reversing the original choice: accepting a gift card framed as a reward, which converts the moment from a loss into something approximating a gain. Each of these is psychologically distinct, and each carries a different ethical and behavioural signature.

Harmon-Jones and Mills (1999) add a detail that matters in a designed choice environment. Dissonance-reduction depends on the consumer's perceived freedom of choice. When people feel they were forced into a behaviour they would not otherwise have chosen, the dissonance is attributed externally and the consumer is less motivated to resolve it internally. The consequence for an interface like Returnflows is that the same architectural feature, a refund fee or a 10% gift card bonus, can land differently depending on whether the consumer experiences the choice as theirs or as the platform's.

### ***2.2.2 Autonomy and self-determination in designed environments***

Self-determination theory (SDT) provides the vocabulary for answering that question. Ryan and Deci (2000) identify three innate psychological needs whose satisfaction predicts the quality of a person's engagement with any activity: competence, autonomy, and relatedness. When these needs are met, motivation is intrinsic or fully internalised. When they are thwarted, motivation becomes controlled, alienated, or anxious, and behaviour suffers even if outcomes are nominally the same.

The critical point concerns what autonomy actually means. Ryan and Deci (2000) define autonomy "not to being independent, detached, or selfish but rather to the feeling of volition that can accompany any act" (p. 73). This is a counterintuitive reading. A consumer who follows a designed flow and chooses the outcome the designer wants can still be autonomous in SDT's sense, provided the choice feels volitional rather than coerced. Conversely, a consumer who technically has three options but feels pressed toward one of them experiences a loss of autonomy even when no option has been removed.

SDT proposes a continuum of regulation running from external to integrated. External regulation is the purest form of controlled motivation, driven by immediate contingencies such as rewards and penalties. Introjected regulation is slightly more internalised but still driven by guilt or self-esteem concerns. Identified regulation occurs when the person consciously values the behaviour. Integrated regulation is fully autonomous, aligning the

behaviour with the person's broader self-concept. Ryan and Deci (2000) are clear that autonomy-supportive conditions, such as a meaningful rationale, acknowledgment of the person's perspective, and real choice, move regulation toward identification and integration. Threats, deadlines, and imposed goals move it the other way.

For Returnflows, this framework poses a specific diagnostic question. If the interface provides a clear and accurate rationale for exchange (faster delivery, avoided fee), the consumer can endorse the choice on identified grounds, and the autonomy criterion is met. If the interface conceals consequences or frames the alternative as a punishment, the same behavioural outcome can occur under introjected regulation, where the consumer chooses what is presented but feels controlled. The platform data does not directly distinguish these states, but the literature predicts behavioural traces. Identified regulation tends to be stable and resilient. Introjected regulation is more fragile and more likely to produce later complaints, refund disputes, or churn. Section 5.2 returns to this distinction.

### ***2.2.3 Moral licensing and the downstream effect of virtuous choice***

The third mechanism matters because the consequences of a return decision do not end when the return is processed. Blanken et al. (2015) define moral licensing as the effect by which past good deeds "can liberate individuals to engage in behaviors that are immoral, unethical, or otherwise problematic" (p. 540, citing Merritt et al., 2010). The phenomenon has been demonstrated across domains. A consumer who purchases eco-friendly products later behaves less ethically in unrelated tasks. A person who recalls a charitable act becomes more willing to act selfishly afterwards. The good deed does not have to match the domain of the subsequent lapse. It licenses across categories.

The mechanism operates as a credit on subsequent self-perception, raising the threshold at which an indulgent or careless act registers as morally problematic. The relevance to a system like Returnflows lies in the framing of the gift card option, examined empirically in section 5.3.

For a return flow that directs consumers away from refund and toward exchange or store credit, the question is whether the diverted choice is experienced as morally meaningful. If it

is, moral licensing predicts that the same consumer may be more permissive in subsequent purchase decisions, more likely to add an item to the basket, less likely to scrutinise the next size choice. The relevant empirical signal in the Returnflows data is the gift card redemption pattern, examined in section 5.3, which provides a partial test of whether store credit produces extended consumption or simply deferred substitution.

Whether a return decision is experienced as a moral act is empirically open. Some consumers may treat the choice as purely pragmatic, in which case the licensing mechanism never engages. Others may process it as virtuous, in which case it does. The thesis returns to this question in the sustainability discussion, where the observed behaviour from the Returnflows dataset provides a partial test of which framing dominates.

*Taken together, these three mechanisms describe the psychological arc of the return decision. Dissonance initiates it. Autonomy, or the consumer's experience of it within the designed flow, determines whether the resulting choice is genuinely owned. Moral licensing describes what may happen afterwards, when the consumer carries the memory of a virtuous choice into the next purchase. But individual psychology is only half the picture. The aggregate consequences of many such decisions, for the environment and for the circular economy more broadly, are the subject of the next section.*

## **2.3 Sustainability, circular economy, and the rebound effect**

### **2.3.1 Rebound in the circular economy**

The circular economy is usually sold as an environmental good. Recycling, refurbishment, resale, and repair are presented as ways to extract more value from existing material and reduce the need for virgin production. The framework treats consumption as something that can be slowed and looped, with environmental benefit accruing through extended product lives and substitution of secondary for primary materials. Zink and Geyer (2017) trouble this narrative. Their core argument is that there is nothing intrinsically green about circular economy activities, because the same mechanisms that make them economically attractive, lower input costs, recovered value, willingness to pay for sustainability, can rebound by enabling additional consumption rather than displacing it.

Zink and Geyer (2017) identify two mechanisms through which rebound arises. The first is insufficient substitutability: when secondary goods fail to perfectly replace primary ones, they add to total consumption rather than displacing virgin production. The second is a price effect: when circular activity expands supply, prices fall, and the resulting income and substitution effects draw in new consumption. The authors conclude that “simply encouraging private firms to find profitable opportunities in the circular economy is very likely to result in rebound” (p. 600). Profit-maximisation and environmental benefit do not automatically align.

Yerushalmi and Saha (2025) extend the warning with a sector-specific quantification. Using a computable general equilibrium model of the textiles and clothing sector, they find that circular-economy innovations produce a 155% rebound backfire, meaning the intervention drives total production and consumption 55% above their pre-innovation levels. The result holds across innovation intensities. Only a Pigouvian tax of at least 1.25% on textile and clothing production appears to offset the effect, and the authors conclude that meaningful sustainability in the sector will require degrowth-informed policies that explicitly target reductions in production and consumption rather than efficiency gains alone. The claim is uncomfortable, but its empirical basis is hard to dismiss.

For this thesis, two aspects of the Zink-Yerushalmi literature need emphasising. Both analyses operate on the production side, asking what happens when firms reduce per-unit impact through recycling or fibre-to-fibre recovery. Neither examines the return journey, in which a consumer who has already bought a garment decides what to do with a mismatched purchase. The structural conditions that produce rebound in secondary-production markets may or may not apply at the level of the individual return decision. Whether they do is precisely the empirical question this thesis takes up in its sustainability discussion.

### ***2.3.2 Consumer-side rebound***

The consumer-side literature on rebound is thinner but more directly relevant. Ciecelska et al. (2024) studied 639 Polish Gen Z consumers to measure whether second-hand clothing purchases displace or supplement new ones. The displacement ratio is striking: consumers who buy only new garments purchase 1.86 items per month on average, while those who buy both new and second-hand purchase 2.28. The implied substitution rate is 1:1.23,

meaning every second-hand item displaces only part of a new purchase and the remainder represents additional consumption. Ciechelska et al. attribute the gap to two mechanisms working in tandem. The price effect draws consumers toward items they would not otherwise have bought, with 36% of their sample admitting to purchasing second-hand garments they did not need because prices were low. Moral licensing, in the Blanken et al. (2015) sense, amplifies the pattern: environmentally framed consumption is reinterpreted as permission to consume more rather than differently.

Makov and Meshulam (2025) reach compatible conclusions through a different method. Modelling 11 million second-hand apparel transactions on eBay, they find that re-spending of savings offsets roughly 47% of the expected environmental benefit, and imperfect substitution raises the offset to 55%. Second-hand commerce delivers genuine environmental value, but headline figures used in industry marketing overstate it by roughly half.

Two caveats matter for how this literature applies to Returnflows. First, all of it concerns secondhand markets, where consumers add used garments alongside new ones. Returnflows operates a different mechanism: an exchange replaces an item the consumer has already purchased but cannot use. There is no price saving and no expansion of demand. Several rebound channels are structurally absent. Second, the channels that do remain are tied to moral licensing and store credit. If a consumer experiences the exchange as a virtuous choice, the licensing risk Blanken et al. (2015) describe remains live. If the 10% gift card bonus leads consumers to treat credit as a windfall, it may expand near-term purchase behaviour. The 22.5% of returns that are not size-related is the portion of the data where these mechanisms would be most visible.

### ***2.3.3 Fashion returns as an environmental problem***

The literature on fashion returns, taken as a whole, establishes the scale of what is at stake. Marriott et al. (2025) report that style and fit dominate return reasons, with sizing accounting for between 28.5 and 40.5% depending on category and bracketing contributing a further 14.5 to 17%. The environmental consequences are substantial: 23 million items discarded in the UK in 2022, 750,000 tonnes of CO<sub>2</sub> from that disposal, 350,000 tonnes of returned apparel sent to landfill in 2017, and returns accounting for roughly 25% of the

e-commerce carbon footprint overall. Their framing is blunt: retailers have historically viewed returns as a routine cost and systematically underestimated their environmental weight.

Stocker et al. (2021) provide the clearest consumer-facing data. Surveying 8,393 German online fashion shoppers, they find that 87.1% of returns are driven by fit, validating the 77.5% sizing share documented in this thesis as a category-level pattern rather than a merchant-specific anomaly. They also document that 41.6% of consumers admit to bracketing. The behavioural pattern is well documented but not directly addressed by current return interfaces, which treat every returned garment as equivalent to every other. The implication for this thesis is that fashion returns are dominated by a particular type of mismatch, the variant being wrong rather than the purchase intent being wrong, and that this dominance is what makes the exchange-favouring architecture analytically interesting.

Synthesising across the three streams, the gap this thesis occupies becomes visible at their intersection. Nudging theory tells us how digital interfaces shape behaviour but has not been tested in return flows. Consumer psychology explains why returns happen and how the experience is resolved but has not been connected to designed choice architecture. Sustainability research documents the aggregate consequences of circular activity but operates at the production and secondhand-market levels, not at the return decision. Returnflows sit where all three converge.

## 3. Case Presentation

### 3.1 Returnflows

Returnflows is a Danish software-as-a-service platform founded in 2024 that replaces the default post-purchase return experience on Shopify merchants. The product grew out of a concrete operational observation. Shopify's native tools support refunds, exchanges, and store credit, but the customer-facing self-serve flow is built around the refund. Anything else falls back on the merchant, who has to handle the case manually from the admin panel by contacting the customer, adding exchange items by hand, issuing credit, and tracking the resolution outside any automated workflow. For both sides this is slow and effortful, and refund becomes the resolution most customers actually receive because it is the path of least friction. Returnflows treats that outcome as a design choice rather than a given, and rebuilds the flow around whether the revenue has to leave the store at all.

Pries Jensen describes returns logistics as a sector that has been historically inefficient and inflexible for shop, warehouse and customer alike (A. Pries Jensen, personal communication, April 28, 2026). The platform is positioned as resolving a multi-stakeholder coordination problem rather than optimising a single line item, with commercial advantages for the shop and matching opportunities for the customer.

The platform is embedded in the merchant's storefront as a branded self-service portal and targets fashion and apparel e-commerce, the category where return rates are highest and where the dominant return cause, size and fit mismatch, is structurally solvable through exchange rather than refund. Its market position is narrow by design. It competes not against logistics providers but against the status quo of email-based, manually handled returns, where the practical default falls on refund because that is the only resolution the standard customer-facing flow automates. The strategic claim is that a return is not a lost transaction but an unfinished one, and that the interface at which the customer decides what to do next is the highest-leverage point in the post-purchase journey.

### 3.2 The return flow

The consumer journey begins when the customer initiates a return through a link provided by the merchant, typically from the order confirmation email or account page. The link opens the Returnflows portal, where the customer authenticates with order number and email. They then select which items to return and why, choosing a return reason. This step serves a double function: it resolves the customer's immediate task, and it generates structured data that later feeds back to the merchant as actionable input for size guides and product descriptions.

Once the reasons are logged, the customer is presented with a resolution screen that constitutes the core choice architecture (see Attachment 1). Three outcomes are offered. Exchange appears first: a full-width row with a "Quick exchange" tag, framed as the fastest path to the right product. Gift card sits beneath, often carrying a bonus on the returned value. Refund sits below, framed around a fee deducted from the reimbursed amount. The ordering, framing, and fee structure are the nudges. None of the options is hidden and none is removed, but the relative cost of each (in clicks, time, money) is deliberately asymmetrical.

If the customer chooses exchange as a swap (same product, different size), they remain inside the return flow and the interface surfaces the same product in alternative sizes. If they choose to exchange for another product, they are redirected to the merchant's shop and browse through the normal storefront, but in a return context. The cart is Returnflows' return cart, displaying what is being returned alongside what is being chosen in its place. The experience mirrors ordinary shopping. Where the exchanged product is more expensive, a payment form appears and the customer pays the difference.

Instant Exchange accelerates the exchange experience independently of whether the swap is to a different variant or a different product. When activated, the replacement item is released for dispatch the moment the parcel is handed in at a drop-off point, or when it arrives at the warehouse, depending on configuration. This collapses the wait from roughly eight days under a refund-and-reorder cycle to around two. Gift cards and refunds both release when the return is closed, which happens either automatically after a configured interval following hand-in or arrival at the warehouse, or after the merchant manually

reviews and closes the return. Instant Exchange is the path that decouples release from close, dispatching the replacement item the moment the return parcel is handed in.

The Optimize Refund Button feature illustrates how granular the architectural logic has become. When activated, the feature demotes the refund control from a full row with a “Continue” action to a plain text link beneath the exchange and gift card rows, which retain their button framing (see Attachment 1). The option is preserved within the libertarian-paternalist frame Thaler and Sunstein (2008) describe, but its perceptual weight is reduced.

The flow closes with confirmation, a shipping label, and a data event written back to the merchant's dashboard. The return is structured as a loop rather than a transaction: the customer resolves a specific dissonance, the merchant retains revenue where possible, and an undifferentiated refund becomes a stream of behavioural data the merchant can act on upstream.

### **3.3 The value proposition**

Returnflows presents itself to prospective merchants in two dimensions, one economic and one environmental, that turn out to be the same dimension seen from different angles. The first is economic. Under the prior manual flow, where almost all returns resolved as refund, the typical return was a one-way transfer of value out of the store. Under Returnflows, a share of that value is redirected into exchanges, which preserve the original transaction in kind, and into gift cards, which convert a cash outflow into deferred in-store credit. The claim is not that returns disappear: return volume is a function of product, sizing and consumer expectation. The claim is that the outcome of each return is a variable rather than a constant, and shifting the distribution even modestly across thousands of returns compounds into material revenue retention.

Pries Jensen frames the work of selling Returnflows to merchants as essentially the work of making the case that good service has a measurable financial return: the proposition is not framed primarily as a sustainability story, nor as a customer-experience story in the abstract, but as a unit-economics story whose plausibility to the merchant is what the platform's

commercial uptake actually depends on (A. Pries Jensen, personal communication, April 28, 2026).

A second order economic effect operates on customer lifetime value (CLV). The return moment is a fragile point in the customer relationship. The purchase has already failed on its original terms, and how the merchant handles the resolution shapes whether the customer returns to the store or churns to a competitor. An exchange that delivers the correct variant, or a gift card that converts a loss into continued engagement, preserves the relationship as well as the revenue.

The second dimension is environmental and operates upstream. Each return reason captured in the portal is written back to the merchant as structured data, and the aggregate constitutes a diagnostic map of where the catalogue misaligns with consumer bodies. A merchant who observes that a specific style is returned disproportionately for being too small can update the size guide, the product description or the fit photography, which reduces future mis-purchases at source. The environmental logic is therefore prevention before diversion: the exchange nudge resolves a return that has already happened, but the data loop attacks the upstream cause.

The two dimensions converge in a single strategic proposition. A return that is converted into an exchange retains revenue and resolves the transaction in a single return and replacement cycle rather than the refund and reorder pattern that typically produces two. A return that is prevented altogether by a corrected size guide removes the shipment, the handling cost and the carbon footprint entirely.

### **3.4 The empirical case**

The merchant studied in this thesis is the Danish direct-to-consumer fashion brand Skagen Clothing, founded in 2021 and headquartered in Aarhus, operating on Shopify and selling primarily to Denmark, Germany and Sweden. Raw transaction-level files were not appended following a prior agreement on commercial sensitivity. Platform data is reported at aggregated monthly or total-period level, with the procedure detailed in chapter four.

The case is suitable for this study for three reasons, ordered by their weight in the selection.

First, the merchant is a fashion brand, the category in which size and fit problems dominate fashion returns, placing the case inside the theoretical frame developed in chapter two rather than at its margin. Second, the case offers a clean before-and-after contrast. Before October 2024, the merchant ran the default Shopify return process, where exchange and store credit each required manual handling through the admin panel. In practice the overwhelming majority of returns resolved as refunds, because that was the only resolution the customer could obtain without merchant intervention. The baseline against which Returnflows is measured is therefore not another sophisticated return solution but the absence of one. Third, eighteen months of post-implementation data allows the analysis to work with steady-state behaviour rather than launch-period noise, and the merchant is one of the platform's longest-running fashion customers.

The case-fit is corroborated by the merchant's own account. In a public Returnflows case study (Returnflows, 2024), Skagen Clothing's co-founder Thomas Melgaard describes returns as a structural feature of fashion e-commerce rather than a defect to be eliminated: "We have improved the size guide and many other aspects, but return volume remains an inevitable part of the business" (Melgaard, as cited in Returnflows, 2024). The framing matters because the merchant's strategic stance lands on the terms this thesis examines: returns are treated as a structural feature of the category, and the design question is what to do with the return event once initiated. Returnflows (2024) is the platform's own publication and reflects the platform's narrative purpose. The source is used here for what the merchant says about itself, while the platform-performance figures used in the analysis come from underlying transaction-level data accessed directly.

The key figures previewed here frame the analysis that follows. Across the observation window, the merchant processed 17,938 returns through Returnflows with a total return value of 7.32 million DKK. Under the prior manual flow, the overwhelming majority of return value resolved as refund. Under Returnflows, that share fell to 62.27%, with 30.08% redirected into exchanges and 7.65% into gift cards. The retained ratio reached 47.16%, corresponding to 3.45 million DKK of revenue that, under the prior manual flow, would most likely have left the store as refund. These figures are unpacked analytically in section 5.2.

## 4. Methodology

### 4.1 Philosophy of science

The thesis takes a pragmatist position, in the sense Saunders, Lewis and Thornhill (2023) describe. Pragmatism judges knowledge by whether it explains a real situation usefully, rather than by whether it matches an external truth or fits inside a single interpretive frame. That suits the research question. Returnflows is a working piece of software, and the relevant test is whether established theory explains what happens when consumers meet it, and what follows for the merchant who runs it.

A strict positivism would not accommodate the interpretive work needed to link UI elements to mechanisms like dissonance reduction or perceived autonomy, which cannot be observed directly. A strict interpretivism would not honour the quantitative platform data that supplies the strategic-value argument. Pragmatism allows both, because both serve the question. The standard for a good answer becomes whether the chosen frameworks produce a defensible account of how the flow works, for whom, and with what side effects.

### 4.2 Research design

The design is a single-case study of Returnflows, with the case merchant introduced in chapter three serving as the empirical site within which the platform is observed in operation. Yin (2018) treats the single case as appropriate when the phenomenon is rare, revelatory, or critical for testing a theoretical proposition. Returnflows qualifies on the second of these criteria. Among Shopify-based return management platforms with publicly visible architecture, it is one of the few designed deliberately around a three-way choice with explicit Communication, Business Administration and Social Diagnostics consequences, and the case merchant's data offers the kind of before-and-after contrast that allows the platform's effects to be analysed against a clean baseline. The single case is a deliberate choice rather than a constraint. Yin (2018) treats the single case as appropriate where the case is critical for testing a theoretical proposition, and the methodological strength of the design lies in depth and counterfactual clarity rather than statistical generalisation.

Generalisation is bounded by definition, and section 4.4 addresses the limits of that bound directly.

The analytical approach is abductive. Established theory, specifically digital nudging, cognitive dissonance, self-determination and rebound, supplies the primary frameworks through which the empirical material is read. The deductive moment is genuine: theoretical mechanisms generate predictions about which UI elements should produce which behavioural distributions, and the platform data is examined for evidence consistent or inconsistent with those predictions. The abductive moment is operationalised through the founder interview, which surfaces practitioner observations that modify what theory alone would predict, most visibly in sections 5.2.3 and 5.3.2 where Pries Jensen's framings reshape the autonomy and rebound discussions (A. Pries Jensen, personal communication, April 28, 2026).

### **4.3 Data collection and sources**

Both datasets were accessed in the author's capacity as an employee of Returnflows under a prior agreement with the merchant and the supervisor. As a condition set by the merchant, the data is cited as an internal source throughout the analysis and anonymised by removing merchant-identifying metadata and presenting figures at an aggregated monthly or total-period level.

The third source is a semi-structured interview with Alfred Pries Jensen, co-founder of Returnflows, conducted on 28 April 2026. The interview followed a prepared guide organised around four themes: origin and vision, the design of the choice architecture, data and behavioural patterns, and the sustainability and rebound argument. It was conducted in Danish, recorded, transcribed, and translated into English, and both versions are reproduced in Appendix A. The format follows the semi-structured interview protocol described by Kvale and Brinkmann (2015), with a fixed set of priority questions and room for the informant to develop adjacent points the interviewer had not anticipated.

The founder interview is not used as an independent source of truth about consumer behaviour, which it could not provide. It is used as a source of design rationale,

strategic framing and practitioner observation that complements the quantitative material, and in two specific places (sections 5.2.3 and 5.3.2) it supplies the abductive element of the research design by surfacing observations that modify the predictions theory alone would generate. The author and the founder are close colleagues at Returnflows, a proximity that grants access but carries bias risk, and the interview material is read with that limitation acknowledged in section 4.4. Secondary sources, including the academic literature reviewed in chapter two, the Returnflows public case study (Returnflows, 2024), and industry reports such as the Business of Fashion and McKinsey State of Fashion 2026, supply the comparative baseline and the market context, cited according to APA 7.

#### **4.4 Analysis, quality criteria and limitations**

Analytically, each theoretical mechanism is mapped onto a specific UI element or behavioural outcome, and the empirical data is then examined for evidence consistent or inconsistent with the predicted pattern. The loss-aversion prediction, for example, generates the expectation that imposing a refund fee while keeping exchange free will shift the resolution distribution away from refund, and the platform data provides the distributional evidence against which that prediction is tested. Reliability is supported by the fact that the quantitative figures are drawn from live production systems rather than from self-report or recall, and validity is supported by the before-and-after structure, which provides a counterfactual approach most digital-nudging studies cannot access.

Generalisability is the weakest criterion. A single merchant cannot prove an industry-wide pattern, and the conclusion accordingly frames findings as propositions for further testing rather than as established claims.

Two further limitations deserve explicit statements. First, the thesis has no direct access to consumer voice. Return behaviour is inferred from what consumers do in the interface, not from what they report about their reasoning. This means claims about dissonance reduction and moral licensing are theoretically grounded rather than cross-checked against consumer reports on the consumer side. Second, the rebound question requires long-term purchase data at the consumer level to be answered definitively, and the present data does not

extend to that level. The discussion therefore argues about rebound risk on the basis of return reason composition and gift card usage patterns, which are informative but not decisive.

The insider position of the author, employed at Returnflows, is treated as both an asset and a liability. As an asset, it grants data access and product-level understanding that an external researcher could not obtain. As a liability, it introduces the risk of motivated interpretation. The thesis limits this risk in three ways: by using pre-specified theoretical predictions that can fail against the data, by highlighting disconfirming evidence where it appears, and by reporting the refund fee and its framing as potentially coercive rather than obviously harmless.

In line with CBS guidelines, generative AI was used as a linguistic and structural aid during drafting and for idea generation during the structuring of the thesis. All analytical judgements, theoretical selections and empirical interpretations are the author's own.

## 5. Analysis and Discussion

The analysis proceeds in three movements. Section 5.1 examines the choice architecture of the Returnflows flow and connects each UI element to the theoretical mechanism that predicts its effect. Section 5.2 turns from mechanism to outcome, examining the behavioural distribution it produces and the strategic value it represents for the merchant. Section 5.3 examines the sustainability implications, asking whether the same architecture that retains revenue also risks licensing the overconsumption it claims to mitigate. Throughout, Communication, Business Administration and Social Diagnostics are treated as aspects of the same object rather than as separate layers stacked on top of each other.

### 5.1 The choice architecture: how Returnflows nudges the return decision

#### *5.1.1 Exchange as the default: visual hierarchy and salience*

The first and most consequential design decision is where exchange sits relative to the other options. On the resolution screen, exchange is presented first in reading order, occupies the largest button, carries the merchant's brand colour and sits visually above gift card and refund. Gift card appears as a secondary button, and below the refund button, when the Optimize Refund Button feature is enabled, is a plain text link. Thaler and Sunstein (2008) argue that no choice environment is free of design. The question for the choice architect is therefore not whether to shape the decision but how. Returnflows simply makes that shaping explicit.

The operative mechanism is a fusion of two forces. The first is the default effect. Although no option is pre-selected technically, the visual dominance of exchange produces what Weinmann et al. (2016) classify as a de facto default. The second is salience, which determines which alternative enters active consideration at all. In combination, default and salience do not merely influence choice. They constitute the choice set the consumer perceives in the first place.

What makes this architecture analytically interesting is what it does to the psychology of the moment. A consumer initiating a return has already experienced dissonance in Festinger's (1957) sense, because an expectation about the purchased garment, typically its fit, has

failed to match reality. Refund requires the consumer to abandon the original purchase decision and accept that the purchase was a mistake. Exchange preserves the original decision by correcting its environment, swapping the wrong size for the right one. This is the cheapest dissonance-reduction path available, and the architecture routes the consumer toward it.

The behavioural consequence is visible in the data. The case merchant resolved 30.08% of return value as exchange and 7.65% as gift card, against a baseline where almost all return value left the store as refund, because exchange and store credit each required manual merchant handling that does not scale to the volumes a fashion merchant produces. Of 7.32 million DKK in total return value, 2.76 million DKK was redirected away from refund into outcomes that preserve the original transaction in kind or convert it into deferred in-store credit. The design move is a Communication choice, the psychological pull a Social Diagnostics mechanism, and the retained revenue a Business Administration outcome. None exists in isolation from the others.

The architectural choice that puts into practice this default is more specific than visual hierarchy alone. Pries Jensen identifies the Shop Overlay as the central design move, the feature that lets the customer “actually come into the shop and sort of experience shopping anew” (A. Pries Jensen, personal communication, April 28, 2026). The exchange experience reuses the merchant’s storefront rather than replacing it: the consumer browses the full catalogue from inside the return flow, and the visual cues that produced the original purchase operate again on the replacement choice. This is salience compounded with familiarity, exploiting both the choice-architecture mechanisms catalogued by Weinmann et al. (2016) and the merchant’s pre-existing investment in conversion design.

### ***5.1.2 The refund fee as loss aversion nudge***

Layered on top of the visual hierarchy is a monetary asymmetry: exchange is usually free the first time, but refund carries a fee, set in the case merchant’s configuration at 49 DKK and deducted from the reimbursed amount. On a transactional view this is a modest sum, roughly a tenth of the merchant’s average order value. On a psychological view it is

considerably more. Thaler and Sunstein (2008) emphasise that losses weigh roughly twice as much as equivalent gains in decision-making, the asymmetry Kahneman and Tversky (1979) formalised as loss aversion. A 49 DKK deduction from a refund is not processed as a small cost set against the return value. It is processed as a loss relative to the amount the consumer expected to recover, and in that frame the same 49 DKK becomes feels much worse than the actual amount.

The architectural effect is to reframe the choice between exchange and refund as asymmetric in economic terms as well as visual ones. Exchange is free, fast and prominent. Refund is less prominent and now also costly. The libertarian-paternalist frame Thaler and Sunstein insist on is preserved in the narrow technical sense that refund remains available. But the fee introduces a question 5.1.1 did not force. When the choice is no longer only visually shaped but also monetarily penalised, is the consumer still exercising autonomous regulation in Ryan and Deci's (2000) sense, or has the nudge slid toward controlled regulation driven by the desire to avoid a penalty? Section 5.2.3 examines this directly. The point here is structural: the refund fee does real behavioural work precisely because it activates loss aversion.

### ***5.1.3 The gift card bonus as gain framing***

If the refund fee operates in the loss domain, the gift card bonus operates in the gain domain. The case merchant offers a 10% uplift on returned value when the consumer selects gift card, surfaced on the resolution screen as a concrete figure rather than an abstract percentage. The mechanism corresponds to what Weinmann et al. (2016, Table 1) classify as an incentive nudge combined with gain framing. A consumer returning 500 DKK worth of product is shown not a neutral 500 DKK credit but a 550 DKK credit, and the additional 50 DKK is perceived as a reward rather than as an avoided loss. The psychology of a gain is qualitatively different from the psychology of an avoided loss, even when the underlying economics are similar (Kahneman, 2011).

The founder's account of the gift card economics introduces a complication the choice-architecture frame does not fully capture. Pries Jensen observes that "a high degree of the gift card business lies in that people don't use the gift cards" (A. Pries Jensen, personal communication, April 28, 2026): a meaningful share of issued credit is never

redeemed. This matters for moral licensing in two directions. The licensing mechanism Blanken et al. (2015) describe operates at the moment of redemption: if the credit is never redeemed, the licensing pathway never engages downstream behaviour. But the framing of the bonus as a reward is the same regardless of whether redemption follows, which means the licensing risk is generated at the moment of choice. The 27.34-day average usage time is therefore a composite of two mechanisms with different economics.

The data show that the bonus pulls behaviour, though less dramatically than the exchange nudge. Gift card resolutions account for 7.65% of return value, with 1.05 million DKK flowing through the pathway. The average gift card is redeemed after 27.34 days, most within sixty, suggesting purposeful rather than impulsive deployment. The same gain framing that makes the bonus effective as a nudge also creates the conditions under which the choice could be experienced as a moral gain. If consumers construe the gift card selection as the responsible choice, the bonus may reinforce that reading rather than merely reward it. The architecture therefore carries a second-order risk the fee structure does not: it can be effective exactly because it makes consumers feel good about the decision, and that feeling is itself the mechanism through which licensing would operate.

#### ***5.1.4 Friction design: instant exchange and the cost of refund***

The fourth element of the architecture is about timing. Under the Instant Exchange variant, the replacement item is dispatched the moment the return parcel is handed over at a drop-off point, which collapses the time to a correctly sized garment from roughly eight days under a standard refund-and-reorder cycle to around two. Refund remains subject to the traditional inspection-and-reimbursement sequence. The asymmetry is not the result of refund being deliberately slowed. Refund and gift card both release when the return is closed, which depends on an automatic timer after the parcel reaches its destination or on a manual review by the merchant. Instant Exchange decouples the replacement dispatch from the close cycle entirely, releasing the new item at hand-in. The theoretical signature of the resulting gap is friction. The architecture removes effort from the preferred path while leaving the standard processing time on the less preferred ones.

Thaler and Sunstein (2008) note that friction operates as a forcing function: removing effort from a desirable path and adding effort to an undesirable path reliably redirects behaviour.

Weinmann et al. (2016) translate this into the digital domain through click depth, step count and wait time. The behavioural effect at the case merchant is consistent. Refund accounts for only 6.95% of return outcomes by count, a share incommensurate with what would be expected if the three options competed on equal effort terms. The merchant converts what would otherwise be a full reimbursement into either a sustained or deferred transaction, and the consumer receives the correct garment faster than under the legacy flow.

#### ***5.1.5 Synthesis: the architecture as a system***

The four elements function as a system rather than isolated nudges: visual hierarchy makes exchange cognitively obvious, loss framing makes refund costly, gain framing makes gift card a reward, and differential friction aligns the effort gradient with the intended distribution. Together they close off the cognitive routes by which a consumer might default to refund.

Two reference points locate the architecture. The pre-purchase nudge in von Zahn et al. (2024) reduces returns by 2.6%, and Rauh et al. (2024) show that pre-purchase nudges carry spillover risk: they also suppress purchase motivation. Returnflows operates downstream of both effects. It does not attempt to prevent the return but shapes what happens once the return has been initiated, which means the spillover risk Rauh et al. identify cannot apply: the purchase has already been completed. The architecture is a nudge located at the one point in the customer journey where shaping behaviour carries no risk of suppressing the transaction itself.

## **5.2 Behavioural outcomes: what happens to consumer decisions and merchant value**

### ***5.2.1 From a single endpoint to a three-way split***

The clearest way to state what the architecture does is to compare the distribution of return outcomes before and after its introduction. Before Returnflows, the case merchant ran on the Shopify default, where the customer-facing self-serve flow offered only one endpoint, a refund, with exchange and store credit available only through manual handling by the merchant. After Returnflows, the same pool of returns resolves into three outcomes: 62.27% of return value as refund, 30.08% as exchange, and 7.65% as gift card. The architecture has not reduced the volume of returns. It has reorganised what a return means once initiated.

One qualification matters before the distribution can be read correctly. Shopify's returned-quantity rate of 22.17% and Returnflows' return rate of 18.03% appear to contradict each other but measure different things: Shopify classifies every item that comes back as a return, including exchanges, while Returnflows tracks the consumer's resolution choice. Returns have not become more frequent. They have become more differentiated.

The distribution becomes analytically interesting when read against the return reasons data. Across 17,938 returns, 77.5% were sizing-related (41.1% too small, 23.05% too big, 10.49% too short, 2.86% too long). A further 4.59% cited colour preference and 3.46% a mismatch between product and image, leaving 14.17% in the other category. Stocker, Baier and Brand (2021) report 87.1% of returns in a German fashion sample as fit-related, validating the pattern as category-level. The reason distribution tells us that the overwhelming share of returns at the case merchant arise from a failure of the purchased variant, not the purchased garment. The customer wanted this product. The size was wrong.

The practitioner perspective qualifies this functional reading without overturning it. Pries Jensen reports being "very surprised" at the volume and frequency with which a subset of consumers return, using the platform "as a kind of digital fitting room" (A. Pries Jensen, personal communication, April 28, 2026). The observation splits the consumer base. For the modal consumer, returns are occasional events triggered by a specific fit failure and resolved through exchange. For a tail of the distribution, returns are an ordinary part of the purchase pattern, and the architectural levers that retain value for the modal consumer may also enable the routine cadence of the tail. This split is taken up again in section 5.3.

This is where Festinger (1957) becomes more than theoretical garnish. The consumer who receives a too-small garment experiences dissonance between the expectation that framed the purchase and the reality of the received item. The architecture in section 5.1 offers three resolutions, and the empirical distribution shows that when sizing drives the return, exchange is the path a significant minority of consumers take when made available. The 30.08% of return value flowing into exchange represents consumers choosing the dissonance-reduction path that preserves their original decision rather than reversing it. The Communication architecture creates the possibility, the Social Diagnostics mechanism explains why consumers use it, and the Business Administration consequence is retention.

### ***5.2.2 The strategic value of the redistribution***

Quantifying what the three-way split is worth to the merchant requires a baseline and a comparison. The baseline is the prior manual flow, where almost all of the 7.32 million DKK of return value generated during the observation window would have resolved as refund, because any alternative required hands-on processing by the merchant on a per-return basis. The actual post-implementation result: of the 7.32 million DKK, 2.20 million DKK resolved as exchange value, 1.05 million DKK as gift card value, and a further 71,825 DKK as upsell where consumers exchanging into a differently priced variant selected an item above the original value. Together these sum to 3.45 million DKK in retained revenue, corresponding to a retained ratio of 47.16%. Nearly half of the return value that, under the prior manual flow, would most likely have left the store as refund instead stays inside it.

The figure reframes the unit economics of a return. Under the prior manual flow, where almost all returns resolved as refund, the typical return was a pure cost event. Under the Returnflows architecture, the same return can resolve in three ways. An exchange preserves revenue and consumes one reverse-logistics shipment rather than two. A gift card defers revenue as in-store credit with redemption concentrated in the first sixty days, and the 10% bonus functions as marketing expenditure whose conversion to future purchase is observable. A refund remains a pure cost, but is now a smaller share of the portfolio.

A note on what this figure does not mean. A 47.16% retained ratio is not a 47.16% improvement in profitability, because exchanges carry their own fulfilment cost and gift cards represent deferred rather than realised revenue. The strategic value lies not in the headline number but in the structural shift it represents. The question the Shopify default forces is how much revenue to write off. The question the Returnflows architecture forces is how much of the return value can be preserved, and in what form. In a sector where Business of Fashion and McKinsey (2026) document executives responding to margin pressure by raising prices against consumer resistance, a mechanism that redirects nearly

half of return value away from outright reimbursement is a direct intervention in the unit economics of the category.

### ***5.2.3 The consumer experience: autonomy or manufactured choice?***

The outcomes in 5.2.1 and 5.2.2 raise a question that cannot be answered by the outcomes themselves. A consumer who selects exchange has exercised a choice. But the choice was offered in an architecture calibrated to make exchange the obvious option, refund the costly one and gift card the rewarded one. Is this the kind of freely willed selection Ryan and Deci (2000) describe as autonomously regulated, or is it closer to introjected regulation, where the behaviour is technically chosen but driven by the desire to avoid a negatively framed alternative? The question is unavoidable because the architecture's effectiveness is precisely what makes it suspect.

The honest position is that the architecture produces a spectrum rather than a single answer. A consumer who selects exchange because the right size is genuinely what they wanted is on the identified end of the self-determination continuum. A consumer who selects exchange primarily to avoid the 49 DKK fee is closer to the introjected end. The platform data cannot distinguish between these two consumer types. Ryan and Deci's framework matters here precisely because it refuses the binary of coerced versus free. It supplies the vocabulary to name a degree of design influence that is real without being illegitimate.

The practitioner perspective sharpens this question rather than dissolving it. Asked where the line falls between motivation and manipulation, Pries Jensen draws it explicitly at the design of visibility:

If we start manipulating the customer, it would be by adding some dark patterns that hid some of the options. If we hid the fact that you can get your money back, so you'd have to click on some little button [meaning, diverting the customer to a screen where all options are not visible]. [...] What we do is that we give the customer all the options at once. (A. Pries Jensen, personal communication, April 28, 2026)

The position is defensible in the libertarian-paternalist frame Thaler and Sunstein (2008) set out. Returnflows places refund on the same screen as exchange and gift card, names the fee, names the bonus, and surfaces all three pathways within a single decision step. No option is hidden, and no information about the consequence of each choice is concealed at the moment of selection. By the standard Pries Jensen sets, the architecture is not a dark pattern.

A second observation adds a comparative anchor. Pries Jensen contrasts Returnflows with manipulation patterns common in commerce: free-shipping thresholds that induce additional purchases through psychological framing, and supermarket multi-buy offers that work because of “a big yellow sign” even when per-unit savings are marginal (A. Pries Jensen, personal communication, April 28, 2026). His view is that those patterns cross the line because they are “actually bad business for the end customer,” whereas Returnflows preserves both the option and the cost-of-each-option in plain view. The framing maps onto the self-determination continuum: transparency about consequences enables identified regulation, while hidden consequences would push the consumer toward introjected or external regulation.

Pries Jensen’s account does not settle the question but bounds it. The architecture is not neutral and not coercive in the strict sense. Its ethical character depends on whether the consumer is functionally aligned with the exchange outcome by their original purchase intent (in which case the architecture facilitates a choice they would have made anyway) or one of the residual cases where purchase intent has shifted (in which case the refund fee operates as a soft penalty). The 77.5% sizing share establishes a strong prior that the first case dominates.

### **5.3 The sustainability paradox: rebound risk at the point of return**

The analysis so far has established that the Returnflows architecture redirects nearly half of return value away from outright refund (depending on the category) and that it does so by shaping a psychological moment the prior flow ignored. The remaining question is whether this redirection produces a net environmental benefit or instead joins the long list of circular economy interventions that have delivered less than they promised. The section proceeds in

three moves. 5.3.1 makes the case for net environmental benefit, 5.3.2 develops the rebound counterargument, and 5.3.3 attempts a synthesis that takes both seriously.

### ***5.3.1 The case for net environmental benefit***

The strongest environmental argument for the Returnflows architecture runs through the reason structure of returns. Section 5.2.1 established that 77.5% of returns are sizing-driven: the garment was wanted, but the wrong size was received. The dominant alternative to exchange is not no transaction, because the consumer's intent has not changed. The alternative is refund followed by a new purchase. A refund-and-reorder pattern produces two outbound parcels and two return parcels on expectation, because the same fit uncertainty can produce a second mismatch. An exchange produces one return parcel and one outbound replacement, and the fit uncertainty is resolved in a single iteration.

The logistics consequence aggregates upward. Marriott et al. (2025) estimate returns account for roughly 25% of e-commerce carbon footprint in fashion, with 23 million returned items discarded in the UK in 2017, corresponding to 750,000 tonnes of CO<sub>2</sub>. A mechanism that converts a share of refund events into exchange events does not remove items from reverse logistics but reduces round trips per resolved transaction. When the 30.08% of return value flowing into exchange is read this way, the shipment count for each transaction halves relative to the refund-and-reorder counterfactual, and the carbon footprint follows.

A second channel operates upstream. Reason codes collected at the resolution screen feed back to the merchant as structured data. A merchant who observes that a specific style is returned disproportionately for being too small can update the size guide, the product description or the fit photography, which reduces future mis-purchases at source. Nestler et al. (2021) demonstrate the scale: article-level sizing interventions at Zalando reduced size-related returns by between 3.8 and 6.6%. The environmental logic is prevention before diversion. The gift card usage pattern supports this reading, with average redemption at 27.34 days and most used within sixty.

Pries Jensen is explicit about the gap between current practice and the design's stated potential. He describes return-reason data as "so much valuable data hidden" that merchants are not yet acting on at scale, and identifies a roadmap item where the platform

would “directly pass that data on to the end customer, so they can see oops, there are actually more people who buy this in large or small” (A. Pries Jensen, personal communication, April 28, 2026). The argument extends Nestler et al. (2021): surfacing article-level data directly to the consumer would extend the prevention loop one step toward source. The strongest environmental case for Returnflows lies not in what the platform retains but in what it eventually prevents.

### ***5.3.2 The rebound counterargument***

The rebound literature supplies a harder counterposition. Zink and Geyer (2017) argue that profit-maximising circular activities are structurally prone to rebound. Yerushalmi and Saha (2025) quantify this in the textile sector at a 155% rebound backfire, with only complementary policy measures limiting the effect. A commercially successful return management platform is a plausible candidate for exactly the kind of profitable circular innovation Zink and Geyer warn tends to backfire.

The founder’s account complicates rather than refutes this reading. Pries Jensen is candid that “an exchange can contribute to people buying something that they then anyway had intended to exchange. But it can also contribute to people actually getting it exchanged for something they will use” (A. Pries Jensen, personal communication, April 28, 2026), and acknowledges the question of consumption volume directly: “I generally have a thought that we should consume less because we consume with greater thought” (A. Pries Jensen, personal communication, April 28, 2026). The position is a two-tier claim. For the modal sizing-driven return, the exchange substitutes for a refund-and-reorder cycle the consumer would have undertaken anyway, and rebound is structurally absent. For the residual cases, including the gift card pathway and the “digital fitting room” tail, the rebound channel remains live.

Ciechelska et al. (2024) identify two pathways for consumer-side rebound: a price effect, where cheaper circular goods free disposable income for additional consumption, and moral licensing in the sense Blanken et al. (2015) describe. The Returnflows architecture carries a structural analogue to the second. The gift card bonus is framed as a reward, and any reward-framed choice creates conditions under which the selection might be experienced as

a moral gain. If the consumer construes the gift card as the responsible choice, the 27.34-day redemption window stops being neutral deferral and becomes a period of primed shopping, and the architecture has created extended engagement with the catalogue at exactly the moment licensing would predict the consumer is least inclined to scrutinise what they add to it.

Two further risks deserve naming. First, frictionless exchange may lower the cognitive cost of the initial purchase. A consumer who knows a mis-sized garment can be resolved in two days through Instant Exchange may approach the size decision at checkout with less care, the front-loaded version of Calcabrini et al.'s (2020) observation that free returns correlate with impulse-driven ordering. Second, the 14.17% of returns marked in the other category is unclear by definition and could conceal behaviours from genuine dissatisfaction to bracketing. The retained ratio is calculated identically whether the returned item was bracketed at purchase or not.

### ***5.3.3 Synthesis: where does the evidence land?***

A careful synthesis must respect a structural distinction the rebound literature itself forces. Ciechelska et al. (2024) study consumers who add secondhand garments alongside new ones; the 1:1.23 substitution ratio is net expansion of ownership. Makov and Meshulam (2025) study recommerce markets where secondhand items displace new ones at rates below 1:1, producing indirect rebound. The Returnflows exchange does neither. An exchange replaces a garment that did not fit with one that does. No net item is added to the wardrobe. No price discount frees up income for additional consumption. The demand is bounded. The structural conditions that drive rebound in the Ciechelska and Makov cases are largely absent in the exchange pathway.

The same cannot be said as cleanly for the gift card pathway, where the 10% bonus introduces a genuine economic incentive to extend consumption and the reward framing creates the licensing conditions Blanken et al. describe. The 7.65% of return value flowing into gift cards is where the rebound argument has the most traction, and the honest conclusion is that this portion carries a real if modest rebound risk. The architecture would be strengthened by an explicit acknowledgement that the gift card bonus is the feature most in tension with its environmental framing.

On balance, the evidence supports a qualified net positive. The exchange pathway, which carries the majority of retained value, is structurally distinct from the rebound-prone cases because it is bounded in demand and one-for-one in substitution. The gift card pathway carries some remaining rebound risk that is real but bounded by its share of the portfolio. The upstream data loop attacks the root cause rather than the symptom. The thesis cannot prove that Returnflows produces a net reduction in aggregate fashion consumption, and does not claim to. What it can say is that the architecture occupies a structurally different position in the rebound literature than the circular economy interventions the literature has criticised.

Pries Jensen frames the same conclusion in the practitioner's terms: "ultimately, the very best thing would be that people just never sent back, even though we live off of people doing it" (A. Pries Jensen, personal communication, April 28, 2026). The acknowledgement that the platform's commercial model depends on the failure mode it claims to address is not a contradiction but a recognition of the same paradox the rebound literature describes at the sector level. The platform's most honest path is to keep prevention strictly upstream of retention.

## 6. Conclusion

### 6.1 Answering the research question

The thesis asked how Returnflows uses digital design and nudging to guide consumer behaviour during apparel returns, and what strategic value this offers webshops compared to traditional return handling. The four sub-questions framed the answer through the Communication, Business Administration and Social Diagnostics pillars, and the analysis has addressed each in turn.

On the mechanism question, the analysis identified four interlocking UI elements that constitute a coherent choice architecture. Visual hierarchy elevates exchange to a de facto default, loss framing through a refund fee makes the alternative feel costly, gain framing through a 10% gift card bonus repositions deferred credit as a reward, and differential friction between instant exchange and traditional refund processing distributes effort to align with the intended distribution of outcomes. Each element corresponds to a specific mechanism in the digital nudging literature (Thaler & Sunstein, 2008; Weinmann et al., 2016), and the novelty lies in their systematic combination at a decision point the existing literature has largely ignored.

On the behavioural question, the case merchant's post-implementation distribution, 62.27% refund, 30.08% exchange and 7.65% gift card by return value, documents a shift the prior manual flow could not realistically produce at scale, because anything beyond a refund required the merchant to handle each case by hand. Festinger's (1957) theory of cognitive dissonance explains why consumers take the exchange path when made available: when 77.5% of returns are sizing-driven, exchange resolves the dissonance by correcting the variant rather than reversing the purchase, the psychologically cheapest of Festinger's three reduction paths. The redistribution produced 3.45 million DKK of retained revenue across the observation window, a retained ratio of 47.16% against a baseline where almost all of the 7.32 million DKK would have resolved as refund, reframing the unit economics of a return from what was effectively a pure cost event into a portfolio of outcomes.

On the sustainability question, the evidence supports a qualified net positive. The exchange pathway, which carries the majority of retained value, is structurally distinct from the

rebound-prone cases in the literature because it is bounded in demand and one-for-one in substitution: no net item is added to the consumer's wardrobe and no price discount frees income for additional consumption. The gift card pathway carries some remaining rebound risk through moral licensing (Blanken et al., 2015), bounded by its 7.65% share. The upstream data loop attacks the root cause and, on the evidence from Nestler et al. (2021), can do so at meaningful scale. The thesis cannot claim that Returnflows produces a net reduction in aggregate fashion consumption. What it can claim is that the architecture occupies a structurally different position than the circular economy interventions the literature has criticised.

## **6.2 Limitations**

Four limitations bound the conclusions. First, the study rests on a single case merchant. The choice of one case was deliberate and methodologically grounded in Yin (2018), but the trade-off is that distributional findings cannot be generalised to other merchants without further testing. Second, the platform data captures consumer choices but not consumer reasoning, which means dissonance, autonomy, and licensing arguments are theoretically grounded rather than cross-checked against consumer reports on the consumer side. Third, the rebound question requires long-term consumer-level purchase data to be settled, and that data is not available. Fourth, the author is employed at Returnflows, a position that creates the access this thesis depends on but also a bias risk that section 4.4 names and the analysis attempts to mitigate.

## **6.3 Recommendations and perspectives**

For Returnflows specifically, two implications follow. The gift card pathway is the feature most in tension with the environmental framing, and explicit acknowledgement of that tension would strengthen rather than weaken the proposition: the architecture can defend the exchange pathway robustly on sustainability grounds. Second, the data loop that feeds reason codes back to merchants is the most valuable environmental feature, and product development that makes this loop more actionable would shift the environmental argument from diversion to prevention.

For the wider e-commerce sector, the findings suggest that the post-purchase return flow is a choice architecture in the full Thaler and Sunstein sense, and treating it as a neutral administrative step is itself a design choice with distributional consequences. Extended Producer Responsibility legislation in the European Union, together with the per-unit fines documented in the Business of Fashion and McKinsey 2026 outlook, pushes end-of-life garment costs back onto producers. That sharpens the economic case for mechanisms that retain value at the point of return and generate data that prevents mis-purchases upstream.

Pries Jensen describes contemporary e-commerce as splitting between dominant players that survive on scale and smaller merchants that must constantly redefine their proposition, because customer acquisition cost now exceeds product margin on most orders (A. Pries Jensen, personal communication, April 28, 2026). Return architecture is one of the few infrastructure layers where smaller merchants can credibly operate at the unit economics the largest players take for granted. When retention and prevention become compliance-relevant rather than only conversion-relevant, the asymmetry between merchants who run sophisticated post-purchase architecture and those who run Shopify defaults becomes structural.

For future research, three directions follow. A multi-case comparison across merchants with different product mixes and price points would test whether the distributional pattern is a property of the architecture or of this particular merchant. A consumer-side qualitative study would settle the autonomy question. A long-term consumer-level tracking study would close the rebound question. The threefold contribution offered here is a foundation that such future work can build on rather than a closure of the question.

## Literature

Blanken, I., van de Ven, N., & Zeelenberg, M. (2015). A meta-analytic review of moral licensing. *Personality and Social Psychology Bulletin*, 41(4), 540–558. <https://doi.org/10.1177/0146167215572134>

Business of Fashion & McKinsey & Company. (2026). The state of fashion 2026: When the rules change. <https://www.businessoffashion.com/reports/the-state-of-fashion-industry/>

Calabrini, A., Dodds, N., Edwards, R., Gustavsson, A. V., & Mathieu, L. (2020). “Free returns” – Reducing returns in the online fashion industry [Unpublished coursework paper, PB403 Psychology of Economic Life]. London School of Economics and Political Science. <https://www.lse.ac.uk/PBS/assets/documents/PEL-coursework/Free>Returns-Reducing>Returns-in-the-Online-Fashion-PEL-coursework-2020-2021.pdf>

Ciechelska, A., Matel, A., Poskrobko, T., & Sidorczuk-Pietraszko, E. (2024). Circular economy rebound effect in the context of second-hand clothing consumption. *Economics and Environment*, 87(4), Article 635. <https://doi.org/10.34659/eis.2023.87.4.635>

Diggins, M. A., Chen, C., & Chen, J. (2016). A review: Customer returns in fashion retailing. In T.-M. Choi (Ed.), *Analytical modeling research in fashion business* (pp. 31–48). Springer. [https://doi.org/10.1007/978-981-10-1014-9\\_3](https://doi.org/10.1007/978-981-10-1014-9_3)

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford University Press.

Harmon-Jones, E., & Mills, J. (1999). An introduction to cognitive dissonance theory and an overview of current perspectives on the theory. In E. Harmon-Jones & J. Mills (Eds.), *Cognitive dissonance: Progress on a pivotal theory in social psychology* (pp. 3–21). American Psychological Association. <https://doi.org/10.1037/10318-001>

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291. <https://doi.org/10.2307/1914185>

Kvale, S., & Brinkmann, S. (2015). *InterViews: Learning the craft of qualitative research interviewing* (3rd ed.). SAGE Publications.

Makov, T., & Meshulam, T. (2025). ReCommerce rebound effects: The case of second-hand apparel. Proceedings of the 6th Product Lifetimes and the Environment Conference (PLATE2025). <https://doi.org/10.54337/plate2025-10421>

Marriott, J., Bektaş, T., Leung, E. K. H., & Lyons, A. (2025). The billion-pound question in fashion e-commerce: Investigating the anatomy of returns. *Transportation Research Part E: Logistics and Transportation Review*, 194, Article 103904. <https://doi.org/10.1016/j.tre.2024.103904>

Morvan, C., & O'Connor, A. (2017). *An analysis of Leon Festinger's A theory of cognitive dissonance*. Macat International.

National Retail Federation & Happy Returns. (2025). 2025 retail returns landscape. <https://nrf.com/research/2025-retail-returns-landscape>

Nestler, A., Karessli, N., Hajjar, K., Weffer, R., & Shirvany, R. (2021). SizeFlags: Reducing size and fit related returns in fashion e-commerce. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (pp. 3432–3440). Association for Computing Machinery. <https://doi.org/10.1145/3447548.3467160>

Rauh, C., Straubert, C., & Sucky, E. (2024). Digital nudging to reduce online returns: Can a simple picture save the environment? In AMCIS 2024 Proceedings (Paper 6). Association for Information Systems. [https://aisel.aisnet.org/amcis2024/sig\\_green/sig\\_green/6](https://aisel.aisnet.org/amcis2024/sig_green/sig_green/6)

Returnflows. (2024). Skagen Clothing × Returnflows: Case study [Case study]. <https://www.returnflows.com/en/cases/skagen-clothing>

Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55(1), 68–78. <https://doi.org/10.1037/0003-066X.55.1.68>

Saunders, M. N. K., Lewis, P., & Thornhill, A. (2023). *Research methods for business students* (9th ed.). Pearson.

Stöcker, B., Baier, D., & Brand, B. M. (2021). New insights in online fashion retail returns from a customers' perspective and their dynamics. *Journal of Business Economics*, 91(8), 1149–1187. <https://doi.org/10.1007/s11573-021-01032-1>

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

von Zahn, M., Bauer, K., Mihale-Wilson, C., Jagow, J., Speicher, M., & Hinz, O. (2024). Smart green nudging: Reducing product returns through digital footprints and causal machine learning. *Marketing Science*. Advance online publication.  
<https://doi.org/10.1287/mksc.2022.0393>

Weinmann, M., Schneider, C., & vom Brocke, J. (2016). Digital nudging. *Business & Information Systems Engineering*, 58(6), 433–436.  
<https://doi.org/10.1007/s12599-016-0453-1>

Yerushalmi, E., & Saha, K. (2025). How circular economy innovation can backfire on the environment: Quantifying the rebound effect of the textiles and clothing sector. *Business Strategy and the Environment*, 34(8), 10495–10512. <https://doi.org/10.1002/bse.70135>

Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE Publications.

Zink, T., & Geyer, R. (2017). Circular economy rebound. *Journal of Industrial Ecology*, 21(3), 593–602. <https://doi.org/10.1111/jiec.12545>

## Attachment 1

With Optimized Refund Button

### Select the action you want

Exchange for another product ( DKK 20.00 return label )	Continue →
<a href="#">⚡ Quick exchange</a>	
Gift certificate ( DKK 30.00 return label )	Continue →
<a href="#">⚡ Quick exchange</a> <a href="#">★ 30.00 SEK bonus</a>	

Or [click here](#) to get a refund (10.00 DKK return label)

Without Optimized Refund Button

### Select the action you want

Exchange for another product ( DKK 20.00 return label )	Continue →
<a href="#">⚡ Quick exchange</a>	
Gift certificate ( DKK 30.00 return label )	Continue →
<a href="#">⚡ Quick exchange</a> <a href="#">★ 30.00 SEK bonus</a>	
Refund via payment method ( DKK 10.00 return label )	Continue →

# Management Report

## *Digital Nudging in Fashion E-commerce Returns*

A strategic briefing on the Returnflows choice architecture

Prepared for Returnflows leadership. This briefing condenses an academic thesis on the Returnflows choice architecture into four strategic recommendations, each grounded in the thesis evidence. The architecture itself is presumed familiar; the value of the briefing lies in what an independent analysis adds: a quantified picture of behavioural and revenue outcomes, an explanation of why the design works at the psychological level, and a defensible position on the sustainability question.

### **The problem and the opportunity**

Fashion e-commerce returns are not a cost line that can be squeezed down to zero. They are a structural feature of the sector. Industry data puts the 2025 return volume in the United States alone at 849.9 billion dollars, or 15.8% of retail sales, and the figure climbs to 19.3% for online sales specifically (NRF & Happy Returns, 2025). Fashion sits at the high end of that distribution, with sizing and fit cited as the dominant driver across academic and industry sources. The case merchant data studied in this thesis, covering 17,938 returns across the eighteen-month observation window starting October 2024, shows 77.5% of returns attributed to sizing and fit alone. The problem is not that consumers return clothes. The problem is that they are almost universally handled as refunds, so each such return converts directly into lost revenue and an outbound reverse-logistics trip.

The architecture rests on a single insight, familiar to the leadership team but worth restating in the academic frame the thesis applies: the refund is not the only possible outcome of a return event. A consumer who ordered a jacket in the wrong size does not want their money back in the abstract. They want the right size of the jacket. A consumer who liked what they saw but not what arrived may still want to spend with the store, just not on that specific item. The traditional Shopify return flow, which is where most Danish and European fashion merchants start, automates only the refund path for the customer. Exchange and store credit

exist on the platform but require the merchant to handle each case manually, which does not scale to the volumes a fashion merchant produces. Returnflows replaces that flow with a three-option architecture, and in doing so reframes the return from a terminal event into a branching decision.

### **What the architecture does, and what it has produced**

The Returnflows interface presents three resolution options at the point of return: exchange for a different variant, store credit with a 10% bonus, or traditional refund with a small fee. The design is not neutral. Exchange is visually elevated and surfaced as the de facto default, and processed instantly through the platform. Store credit is framed as a gain, using a 10% bonus on top of the return value. Refund is framed as a loss, using a small handling fee. Friction is distributed so that the option most aligned with retained value is also the option that is easiest to complete.

On the case merchant data, the distribution across these three paths is 62.27% refund, 30.08% exchange, and 7.65% store credit, measured by return value. That distribution is the key operational result. Under the prior manual flow, almost all of the 7.32 million DKK of return value would have resolved as refund, because exchange and store credit each required manual handling that did not scale to the volumes a fashion merchant produces. Under Returnflows, 3.45 million DKK of that value was retained by the merchant, producing a retained ratio of 47.16%. The refund rate, calculated as refunded value over total sales, sits at 6.95%. The unit economics of a return event have shifted from what was effectively a pure cost to a portfolio of outcomes, and the portfolio leans toward retention.

The behavioural mechanism behind this shift is well understood in the literature, even if it has rarely been applied at the return point specifically. The case merchant data shows that 77.5% of returns are sizing or fit related, meaning a typical return is functionally a request to correct the variant rather than to reverse the transaction. The psychologically cheapest way to resolve that mismatch is exchange. The prior flow did not surface that option in the self-serve UI. The new flow makes it the path of least resistance.

### **Strategic recommendations**

Recommendation 1: Lean into the environmental story with precision, not breadth. The exchange pathway, which carries the largest share of retained value, has a defensible environmental argument. It replaces two round-trips of reverse logistics with one, prevents a return from terminating in landfill or resale discount, and is structurally distinct from the rebound-prone interventions the literature warns about: the exchange is bounded in demand and one-for-one in substitution, so no net item is added to the consumer's wardrobe and no price discount frees income for additional consumption. The same data loop attacks the root cause when reason codes feed back into merchant size guides. The store credit pathway is harder to defend unconditionally, because the literature on moral licensing and circular economy rebound shows that perceived sustainability can stimulate additional consumption. Merchant-facing communication should defend the exchange mechanism on its strongest grounds and treat the credit mechanism as a business feature rather than an environmental one. A blanket green claim is weaker than a precise one.

Recommendation 2: Invest in the data loop as the strategic core. The reason-code capture that sits underneath the return flow is, in aggregate, more valuable than any single conversion metric. It tells the merchant which variants fit poorly, which descriptions mislead, and which size charts drift from the actual garment. Work on large-scale data from Zalando shows that article-level sizing interventions reduce size-related returns by between 3.8 and 6.6% (Nestler et al., 2021). The product priority implication is that the analytics side of Returnflows, the part that tells merchants what to change upstream, should be developed with the same seriousness as the return-flow UI itself. Returns prevented are worth more than returns retained.

Recommendation 3: Position Returnflows as regulatory infrastructure, not just conversion tooling. The regulatory direction in Europe is toward Extended Producer Responsibility for textiles, with per-unit fines on unsold and discarded garments becoming a real cost line in 2026 and beyond (Business of Fashion & McKinsey, 2026). A return architecture that retains value at the decision point and generates data that prevents mis-purchases upstream is not only a conversion tool. It is compliance infrastructure. The sales narrative should reflect that. A merchant evaluating Returnflows in 2026 is not choosing between return software and no return software. They are choosing between a return flow that helps them with EPR and one that does not.

Recommendation 4: Treat the gift card tension as a feature of the product, not a bug. The 7.65% share of return value flowing through the gift card path is the smallest pathway, and it is the one most exposed to the moral licensing critique. The 27.34-day average redemption window points to purposeful rather than impulsive deployment, which contains the rebound risk in practice. Rather than defending the pathway as environmentally neutral, the stronger position is to acknowledge it openly as a commercial feature that preserves customer goodwill and reduces churn, and to pair it with strong exchange and data-loop messaging. Merchants and regulators both respond better to a company that names its tradeoffs than one that claims to have none.

### **Closing note**

The thesis this report summarises finds that the Returnflows architecture produces a genuine, measurable redistribution of return outcomes, with a retained ratio near 47% and a structural position on sustainability that is defensible, though not unconditional. The commercial case is straightforward. The environmental case is strongest when stated precisely. The regulatory case is accelerating. The recommendation, in one sentence, is to invest in the data loop, defend the exchange pathway on explicit grounds, acknowledge the gift card tension, and position the product as EPR-era infrastructure. That combination matches what the evidence shows and what the market in 2026 will reward. The findings rest on a single case merchant, so the precise distribution will not transfer unchanged to other merchants, but the architectural logic and the rebound argument do.