

C.A.S. Full Assessment Report

Temu — Temu Ecommerce

Generated: August 27, 2026, 1:56 AM EDT • Assessment ID: 6a8f641c7ed44f83f577c516

Author	Delante Clark
ORCID	0009-0001-9904-0111
School	University of Central Florida
System	Cognitive Attentional Standards (C.A.S.) Platform
System Designer	Delante Clark
System Components	Attention Safeguard Models (A.S.M.), Sociological Impact Assessment (S.I.A.)
Year	2026
Subject	Temu Ecommerce
Mode	Research Mode (Non-Certifying)

Suggested citation

Clark, D. (2026). *Cognitive Attentional Standards Platform (C.A.S.) Assessment Report: Temu. Attention Safeguard Models and Sociological Impact Assessment Output.*

1. Assessment Details

Organization	Temu
System / Product	Temu Ecommerce
System Type	Platform
Primary URL	https://www temu.com/
Assessment Type	Research
Liaison	N/A
Status	Research Only
License Status	Inactive

Scope Description

This assessment evaluates the publicly accessible Temu e-commerce platform, including product discovery, search, recommendations, promotional offers, countdown timers, discounts, coupons, rewards, cart activity, checkout prompts, notification requests, and account-creation flows. It examines how these features influence attention, browsing, purchasing decisions, autonomy, urgency, and the ability to disengage. Findings are based only on publicly observable website behavior and available platform policies. Private user data, payment information, internal systems, proprietary code, and security testing are excluded. This is a research assessment, not a definitive legal determination.

Included Components

Public website pages, search, product recommendations, promotions, discounts, coupons, rewards, countdowns, cart and checkout flows, notification prompts, account prompts, user controls, and exit pathways.

Excluded Components

Private accounts, personal and payment data, order history, proprietary code, internal systems, private communications, mobile-app features, security testing, and definitive legal determinations.

Assessment Contributors

Contributor	Role	Contributions
delante@asmcas.com	Admin	Finalized SIA; saved SPI, AIS, and AEA scores; created assessment

2. Preliminary Assessment Report

Generated August 26, 2026, 6:23 PM EDT · Pages scanned: 200

Executive Summary

The research-based assessment of the Temu e-commerce platform reveals structural challenges regarding accessibility and user autonomy. Automated scanning across 200 pages identified a consistent lack of navigational landmarks, clear exit mechanisms, and instructional content. The platform exhibits behaviors associated with high-pressure sales environments, which may impact user cognitive autonomy. Further manual analysis is required to contextualize these findings against established Cognitive Accessibility Standards.

Detected Signals (5 total)

Category	Signal	Description
AEA	Mobile vs. Desktop Differences	200 pages missing viewport meta tag, potentially degrading mobile experience.
SPI	Missing Clear Instructions	200 of 200 pages lack visible help or instructional content.
SPI	Unpredictable Navigation	200 of 200 pages lack a semantic <nav> element or navigation landmark.
SPI	No Clear Exit or Pause Options	200 of 200 pages do not appear to offer exit, cancel, or pause functionality.
SPI	No Time Expectations Communicated	Time-related cues (e.g., estimated read time) not found on 200 of 200 pages.

Next Steps

- Initiate a manual Subject Matter Expert (SME) audit to evaluate user experience flows, particularly high-pressure purchase prompts and navigational clarity.
- Proceed to conduct a full Sociological Impact Assessment (SIA) to determine the impact of these identified signals on user autonomy.

3. C.A.S. Scoring (AEA / AIS / SPI)

Metric	Score	Band	Assessor	Date
AEA — Attention Equity Assessment	19.5 / 100	High Inequity	delante@asmcas.com	Aug 26, 2026
AIS — Attention Inequity Score	38.3 / 100	Low Inequity (near threshold)	delante@asmcas.com	Aug 26, 2026
SPI — Safeguard Pressure Index	38.7 / 100	Low Pressure (near threshold)	delante@asmcas.com	Aug 26, 2026

Highest burden group: Mobile-only, lower-resource users relying on smartphones for shopping access.

Scores Summary

Temu's CAS scoring profile is mixed but unfavorable overall. The AEA score of 19.5/100 is the most consequential result because AEA is a higher-is-better metric; this score falls in the High Inequity range and indicates very low attention equity. The AIS score of 38.3/100 and SPI score of 38.7/100 both remain in their respective Low Inequity / Low Pressure bands, since those are higher-is-worse metrics, but both are close to the moderate-risk threshold of 41, above which they would receive a worse classification. In summary, the assessed public website did not reach CAS's more severe formal classifications for inequity extraction or coercive pressure, but it performed poorly on the presence of equity-supportive design features and showed multiple near-threshold pressure indicators. This is not a strong compliance posture.

4. Sociological Impact Assessment (SIA)

SIA Mode	Mode A — Include SIA
Status	Complete
Assessor	delante@asmcas.com
Finalized	Aug 26, 2026

Age Impact Cohort — AIO (2 triggers)

[AIO-A2] Design: Countdown-based promotional offers, expiring coupons, and urgency-framed checkout prompts. Harm: Hurried purchasing or unavoidable loss of promotional eligibility when users cannot respond within imposed time windows.

[AIO-O1] Design: High-density promotional interface with layered offers, multi-step shopping flows, and limited visible instructional support. Harm: Navigation errors, transaction mistakes, and exclusion from full use of the platform's offer and checkout flows.

Digital Divide — DDIO (4 flags)

[DD-1A] Harm: Reduced or blocked access to the platform's core functions and promotional benefits.

[DD-2A] Harm: Errors, task failure, and process abandonment during product selection and purchase.

[DD-3A] Harm: Disproportionate loss of discounts or price access from delayed or interrupted completion.

[DD-3B] Harm: Unequal benefit distribution despite equal formal access to the platform.

Targeted Remediation Entries — TRE (6 required institutional fixes)

hpa-aio-AIO-A2

- **L1 (Design):** Remove or substantially soften countdown-based urgency in purchase-critical flows and replace expiring prompts with stable availability messaging and clear hold indicators.
- **L2 (Structural):** Redesign promotional architecture so offers remain valid through an asynchronous completion window rather than requiring immediate action during browsing or checkout.
- **L3 (Governance):** Adopt a formal promotional-timing standard that prohibits strict deadline design for core purchase decisions without documented grace periods and reviewable exception handling.

hpa-aio-AIO-O1

- **L1 (Design):** Reduce interface density, consolidate overlapping promotional modules, and add persistent step cues, visible help text, and predictable navigation labels.
- **L2 (Structural):** Create simplified browsing and checkout pathways with saved state, persistent cart and offer context, and minimized cross-page recall requirements.
- **L3 (Governance):** Require cognitive-load and age-access review for promotional and checkout releases, with sign-off criteria tied to instruction visibility and navigation predictability.

hpa-ddio-DD-1A

- **L1 (Design):** Provide a lightweight, low-data web path with visible support contact options and a no-account browsing path for core discovery and cart access.
- **L2 (Structural):** Establish an equivalent non-digital or human-assisted order-support channel for core purchasing and issue resolution.
- **L3 (Governance):** Adopt an access-equivalency policy requiring non-digital alternatives or documented parity measures whenever core functions are delivered online only.

hpa-ddio-DD-2A

- **L1 (Design):** Add visible step-by-step instructions, auto-apply eligible coupons, and present plain-language explanations for rewards, discounts, and checkout states.
- **L2 (Structural):** Redesign promotional and checkout workflows into guided, fewer-step sequences with persistent state and reduced branching logic.
- **L3 (Governance):** Require digital-literacy usability review and evidence-based guidance standards before deploying discount, rewards, or account-conversion flows.

hpa-ddio-DD-3A

- **L1 (Design):** Add clear warning states before expiry, preserve visible fallback pricing, and provide immediate recovery prompts when a timed offer lapses mid-flow.
- **L2 (Structural):** Implement grace-period logic that preserves promotional eligibility for interrupted sessions and delayed checkout completion.
- **L3 (Governance):** Create policy rules requiring recovery windows, override authority, and audit review for any promotional mechanism that withdraws value after a missed action.

hpa-ddio-DD-3B

- **L1 (Design):** Simplify price presentation and automatically apply the best eligible discount instead of requiring users to collect, compare, and stack offers manually.

- **L2 (Structural):** Redesign rewards and promotions so value does not depend on repeated monitoring, rapid response, or multi-step optimization behavior.
- **L3 (Governance):** Institute periodic resource-equity audits for pricing, rewards, and promotion systems to ensure benefits are not conditioned on excess time, attention, or digital fluency.

SIA Narrative

SIA status is complete. Mode A was included and completed on August 26, 2026. There were 0 AER exemption records, 0 RSIO flags, and 0 HPA Harm Ledger entries. The SIA is appended to the CAS report and does not affect certification tier, but it provides substantial context for real-world impact.

Racialized Structural Risk (RSIO): 0 flags were identified in this assessment. This should not be interpreted as a universal absence of structural risk; it means no RSIO flags were identified within the scope and method used here.

TRE remediation output (6 findings) focuses on design simplification, timing safeguards, access equivalency, and guided workflows. Recommended actions include softening or removing countdown-based urgency in purchase-critical flows; preserving offer validity through asynchronous completion windows and grace periods; reducing interface density and consolidating overlapping promotional modules; adding persistent step cues, visible help text, plain-language explanations, and predictable navigation labels; creating simplified browsing and checkout paths with saved state and reduced recall burden; providing low-data and no-account pathways with equivalent human-assisted or non-digital support; automatically applying the best eligible discount rather than requiring manual offer optimization; and adopting formal review standards for timing, age accessibility, digital-literacy usability, and resource-equity auditing.

5. Attention Safeguard Model (ASM) Scoring

Source: ASM Scoring Dashboard — platform pressure inputs and simulation scores linked to this assessment. Four completed ASM runs were provided for Temu under different run configurations and target labels.

Run / Target	ASM Score	Status	Recovery	Cycle	Behaviors	Scored At
Existing Platform Input — Temu	84.1 / 100	Meets Standards	5.2 / 10	5	Selective exposure, muting, avoidance, disengagement	Aug 27, 1:29 AM
Temu (platform)	84.1 / 100	Meets Standards	5.2 / 10	5	Selective exposure, muting, avoidance, disengagement	Aug 27, 1:29 AM
Temu (target: Website)	12.0 / 100	Does Not Meet Standards	9.3 / 10	5	Selective exposure	Aug 27, 1:19 AM
Temu (target: Website, earlier run)	12.0 / 100	Does Not Meet Standards	9.3 / 10	5	Selective exposure	Aug 26, 7:06 PM
Temu Ecommerce (target: Website)	40.5 / 100	Partially Meets Standards	8.2 / 10	5	Selective exposure, muting	Aug 26, 7:06 PM

Linked Social Pressure Inputs

Source	API Pressure Score
Temu (assessment import)	75.0 / 100
Website — Temu Ecommerce	35.9 / 100

ASM Narrative

ASM scorer data is supplemental simulation evidence and is not a standalone compliance determination. The ASM outputs vary significantly by run configuration and target labeling, ranging from 12 to 84.1, which suggests the simulation is sensitive to context, scenario framing, or model input path.

Even with that variance, one annotation is consistent across all runs: selective exposure, indicating a repeated simulated need to filter, narrow, or defend attention against competing inputs. Muting appears in three of the four runs, and avoidance and disengagement appear in the two highest-scoring Temu-target runs, aligning with the CAS/SIA picture of a commerce environment that may require users to actively suppress or work around attention-demanding elements.

Recovery scores are low to modest across runs (5.2, 5.2, 8.2, 9.3), suggesting limited restorative margin in at least some simulated interaction states. While the ASM results are not uniform enough to serve as a sole basis for a finding, they reinforce the report's central conclusion that Temu's observable interface can create recurring attentional management demands, especially around selective exposure and disengagement behaviors.

6. C.A.S. Compliance Report — Detailed Findings

Summary

This CAS research assessment examines the publicly accessible Temu e-commerce platform at <https://www.temu.com/>, with emphasis on product discovery, search, recommendations, promotions, discounts, coupons, rewards, countdowns, cart and checkout prompts, notification requests, and account-creation flows. Assessment status is Research Only; this is a non-certifying research assessment and the system is not certified. The review was conducted as a research assessment by delante@asmcas.com and is expressly limited to publicly observable platform behavior rather than internal systems or a definitive legal determination.

The scored results show a mixed but concerning accessibility picture. Temu received an AEA (Attention Equity Assessment) score of 19.5/100; because AEA is a metric where higher is better, this falls in the High Inequity range (≤ 40) and indicates very low attention equity in the observed design. By comparison, the AIS (Attention Inequity Score) was 38.3/100 and the SPI (Safeguard Pressure Index) was 38.7/100. For both AIS and SPI, higher is worse; both scores remain within the Low Inequity / Low Pressure band (≤ 40), but both are close enough to the moderate-risk threshold that a small increase would produce a worse classification. Taken together, this means the platform did not score as severely coercive on the measured inequity and pressure scales, yet it showed a substantial lack of equity-supportive design features that help users orient, proceed deliberately, and disengage safely.

The preliminary crawl supports that interpretation. Across 200 scanned pages, the automated review identified structural issues including a lack of navigational landmarks, limited instructional content, no clear exit or pause options, unpredictable navigation, and no communicated time expectations. The preliminary report also recorded one AEA signal (mobile vs. desktop differences) and four SPI-related signals. No preliminary AIS signals were recorded in the automated scan, suggesting the most salient early concern was not extreme overt manipulation but rather an environment with weak cognitive supports.

The completed Sociological Impact Assessment (SIA), Mode A, materially strengthens the concern profile. It triggered 2 Age Impact Cohort items and flagged 4 Digital Divide items, with no RSIO flags and no AER exemptions. Six targeted remediation entries were issued through TRE, indicating that meaningful redesign would be required to improve autonomy, guidance, pricing recovery, and access equivalency.

Scope Limits

Scope assessed: This research assessment evaluated the publicly accessible Temu website experience, including public website pages, product discovery, search, product recommendations, promotions, discounts, coupons, rewards, countdowns, cart activity, checkout flows, notification prompts, account prompts, user controls, and exit pathways.

Scope excluded: Private accounts, personal data, payment data, order history, proprietary code, internal systems, private communications, mobile-app-only features, security testing, and any definitive legal determination were excluded. Findings are based only on publicly observable website behavior and available platform policies during the assessment window.

Assessment limitations: Because this is a research assessment, it should be read as an evidence-based compliance analysis of observed design conditions, not as a formal certification, enforcement order, or binding legal finding. Status remains Research Only, and certification has not been granted.

Finding 1 — Attention Equity Assessment (AEA): 19.5 / 100, High Inequity

AEA is a positive metric where higher scores are better. Temu's score of 19.5 is substantially below the 40-point threshold and indicates poor attention equity in the assessed experience. In practical terms, the publicly observable design appears to place a high cognitive burden on users while providing limited structural support for orientation, pacing, and informed disengagement.

Key factors associated with the low AEA result:

- High-density promotional presentation, with layered offers, discounts, rewards, and recommendation surfaces competing for user attention.
- Multi-step shopping and checkout flows with limited visible instructional support.
- Countdown-based promotional offers and expiring coupon states that require users to monitor timing and status rather than rely on stable access conditions.
- Limited evidence of clear pause, exit, or recovery pathways once users enter promotionally dense or purchase-adjacent flows.
- Preliminary crawl indicators showing missing navigational landmarks, missing clear instructions, and no communicated time expectations.
- An AEA signal noting mobile vs. desktop differences, suggesting attention support may not be consistent across access contexts.

Interpretation: The AEA score is the clearest indicator in this assessment that the design does not adequately distribute attention demands in an equitable way. Users who need clarity, extra time, lower density, or stable cues are more likely to experience friction, confusion, or rushed decision-making.

Finding 2 — Attention Inequity Score (AIS): 38.3 / 100, Low Inequity band, near threshold

AIS is a negative metric where higher scores are worse. Temu's AIS score of 38.3 remains within the Low Inequity band (≤ 40), which is comparatively better than a moderate or significant inequity classification, but it sits only a few points below the cutoff for Moderate Inequity, so it should not be read as a strong margin of safety.

What this means in context:

- The assessment did not place Temu in the moderate or significant inequity range on this specific scale.
- At the same time, the platform shows recurring patterns that can distribute burden unevenly: promotional complexity, unstable discount states, and success pathways that favor users with more time, digital fluency, or ability to monitor expiring opportunities.
- The SIA DDIO findings reinforce this interpretation by identifying resource-dependent success pathways and disproportionate loss of value when users cannot complete tasks quickly.

Important interpretive note: A low-band AIS score does not override the very low AEA score; the two metrics measure different constructs. Here, the combined result suggests the platform's main issue is not solely extreme overt inequity extraction, but a substantial absence of equity-supportive design that leaves users to manage complexity, urgency, and promotional instability with insufficient support.

Finding 3 — Safeguard Pressure Index (SPI): 38.7 / 100, Low Pressure band, near threshold

SPI is also a negative metric where higher scores are worse; a score of 0 would be the best possible result and would mean no coercive pressure was detected. Temu's SPI score of 38.7 remains in the Low Pressure band, but it is only a small margin below Moderate Pressure. This means the platform did not cross into a formally moderate coercion classification, yet observable pressure-related features are present and meaningful.

Pressure-related concerns identified across the assessment:

- Countdown-based promotional offers.
- Expiring coupons and urgency-framed checkout prompts.
- No clear exit or pause options identified in the preliminary scan.
- Unpredictable navigation and missing clear instructions that can make it harder to slow down or verify choices before acting.
- No time expectations communicated, which can intensify pressure when discounts or offers appear unstable.
- Notification and account prompts within a commerce setting where the user's attention is already heavily contested.

Interpretation: The SPI result indicates that the site's pressure profile is not among the most severe outcomes measurable by CAS, but it also does not indicate the presence of robust safeguards. The score is best understood as a warning-zone low-pressure result: technically below the moderate cutoff, but close enough that persistent urgency mechanics and weak user controls materially affect autonomy.

Overall Finding Across Domains

Temu's risk profile is driven most strongly by weak attention equity and by design conditions that demand sustained vigilance from users. The score pattern — AEA 19.5, AIS 38.3, SPI 38.7 — suggests a platform where overt coercion is not scored at the highest levels, yet the surrounding architecture still makes fair, calm, and recoverable shopping more difficult than it should be.

7. C.A.S. Improvement Roadmap

Temu's highest-priority CAS improvement area is attention equity. The platform's AEA score of 19.5 indicates high inequity in how attention is demanded and managed, and the assessment findings point to the main causes: dense promotional layering, countdown-based urgency, limited instructional support, inconsistent navigation cues, and weak disengagement pathways. Across the 200-page research scan and SIA review, the same pattern appears repeatedly: users are asked to process overlapping discounts, rewards, prompts, and time-sensitive offers without enough stability, explanation, or recovery support. This raises the likelihood of hurried decisions, navigation errors, and unequal success rates for users with lower digital fluency, older users, or users who cannot respond quickly to expiring states.

AEA — Attention Exposure

1. Reduce above-the-fold promotional density on homepage, search, and product pages by limiting concurrent banners, flash-sale modules, coupon prompts, rewards callouts, and countdown elements to one primary offer zone per screen.
2. Create a stable 'standard shopping view' that users can switch to from homepage, search results, product detail, cart, and checkout, hiding nonessential promotional overlays and preserving only core product, price, shipping, cart, and checkout information.
3. Standardize page structure across desktop and mobile so core actions appear in the same order and location: search, filters, product details, price, coupon eligibility, add to cart, cart review, and checkout.
4. Add persistent instructional microcopy near pricing and purchase controls explaining, in plain language, what price is final, whether a coupon is already included, and whether any action is optional.
5. Introduce visible navigational landmarks and a persistent step indicator for cart and checkout flows, including 'Cart,' 'Shipping,' 'Payment,' 'Review,' and 'Place Order.'
6. Replace rotating or rapidly changing promotional modules with static, user-invoked panels that reveal detail only when expanded.
7. Add a persistent 'Pause promotions' or 'Show fewer offers' control that temporarily reduces attention-demanding modules for the session.

AIS — Attention Inequity

1. Preserve Temu's currently low AIS risk position by establishing a release rule that no new promotion, coupon, rewards, or checkout experiment can add another decision branch unless an existing branch is removed.
2. Automatically apply the best eligible discount at cart and checkout instead of requiring users to discover, collect, compare, or stack offers manually.
3. Provide a visible 'Why am I seeing this price?' explanation next to discounted prices, coupons, reward offsets, and urgency-framed offers.
4. Ensure full product discovery and cart-building can be completed without account creation, app installation pressure, or notification opt-in.
5. Add a persistent cart and offer context layer so users returning after interruption can see original price, current price, coupon status, reward status, and expiry status.
6. Introduce plain-language decision support for offers, such as 'Best price already applied,' 'No action needed,' and 'You can check out without joining this program.'

7. Run segmented usability tests with older adults and lower-digital-fluency users before releasing major pricing or rewards changes, with pass/fail criteria tied to coupon comprehension and accurate identification of final payable price.

SPI — Safeguard & Protection

1. Remove or substantially soften countdown-based urgency in purchase-critical flows, especially where expiring coupons or checkout prompts imply immediate loss.
2. Add clear time expectations anywhere a timed state remains: when an offer started, when it ends, whether it can be recovered, and what price remains if it expires.
3. Implement a grace-period policy for timed promotions that protects users interrupted mid-flow, preserving eligibility for a fixed recovery window.
4. Create a visible 'Exit promotional flow' path from search, product detail, cart, and checkout pages that lets users continue with standard checkout without losing core cart progress.
5. Stop presenting urgency as the default framing for core purchase decisions; replace scarcity or immediate-loss framing with stable messages such as 'available now' or 'standard price available if you return later.'
6. Require a safeguard review for all sale-event launches, including countdowns, stock alerts, coupon claim friction, checkout nudges, and re-prompt frequency.
7. Add one-click suppression for notification prompts and repeated offer modals during the session.

SIA — Sociological Impact

1. For AIO-A2, redesign countdown-based promotional offers so users can complete purchases asynchronously within a documented hold window rather than losing access to a core benefit immediately upon delay.
2. For AIO-O1, consolidate overlapping rewards, coupons, and discount modules into a single offer summary component with expandable details, and add visible step cues and plain-language help throughout selection, cart review, and checkout.
3. For DD-1A, provide a lightweight low-data web path for browsing, cart management, and checkout, plus visible customer-support contact options for users who need human assistance.
4. For DD-2A, convert promotional and checkout flows into guided sequences with fewer branches and less recall burden; auto-apply eligible coupons and keep the final payable total visible throughout.
5. For DD-3A, implement recovery logic for interrupted or expired promotional states, including pre-expiry warnings, visible fallback pricing, and immediate recovery options when a timed offer lapses mid-session.
6. For DD-3B, redesign pricing and rewards so the best available value does not depend on repeated monitoring or multi-step optimization behavior; auto-select and explain the best eligible discount.
7. Adopt a formal promotional-timing and access-equivalency policy prohibiting strict deadline mechanics in core purchasing without grace periods, and assign override authority for restoring lost promotions in edge cases.
8. Institute periodic resource-equity audits reviewing whether best-price access is disproportionately dependent on time, attention, connection quality, device performance, or digital fluency.

ASM — Attention Safeguard Model

1. Address the repeated selective-exposure annotation by adding stronger user-controlled filtering and recommendation reset tools, without requiring account creation or deep settings changes.
2. Address the repeated muting annotation by allowing users to silence promotional prompts, popups, coupon reminders, and notification requests for the session or longer, with the control persisting across page navigation.
3. Address the avoidance and disengagement annotations by building frictionless exit pathways from high-stimulation shopping loops, such as 'continue without offers' and 'save cart and leave.'
4. Improve recovery performance by preserving cart state, viewed-item context, and applicable offer context when users step away and return later.
5. Offer a post-disengagement re-entry page summarizing what changed since the user left: item availability, standard price, preserved discount status, and next step.
6. Reduce recommendation saturation during core tasks by limiting recommendation modules during cart review and checkout to one low-salience section or removing them entirely.
7. Add a 'shopping focus mode' that minimizes selective-exposure pressure by showing fewer recommendation rails, fewer urgency prompts, and no rotating sale modules during product comparison and checkout.

This roadmap is an accessibility and user-autonomy improvement document based on a research assessment of publicly observable platform behavior. It is not legal advice, does not determine liability, and should not be used to justify manipulative optimization under the appearance of compliance. Recommendations should be implemented in ways that increase user understanding, reduce coercion, preserve equitable access, and respect user agency across age, literacy, and resource differences.

8. Algorithm Detection Report

Overall Risk Score

78 / 100 HIGH

Temu (temu.com) exhibits a high concentration of algorithmically driven manipulation and dark-pattern mechanics across all 15 crawled pages. Structural signal detection confirms the simultaneous deployment of infinite scroll, countdown timers, variable reward loops, intermittent reinforcement, attention capture mechanisms, forced continuity, manipulative defaults, and roach motel patterns, active across the full breadth of the site's international footprint.

The SIA sociological impact score of 66/100 with 4 confirmed flags corroborates a platform architecture that is systematically designed to maximize dwell time, compulsive purchasing, and user retention at the expense of autonomous decision-making. Temu's model — ultra-low pricing, gamified discovery, and aggressive re-engagement mechanics — constitutes a high-risk attention-exploitation environment, particularly for vulnerable populations including minors and individuals with impulse-control vulnerabilities. No urgency language or hidden costs were detected in the crawled text sample; however, the structural patterns alone represent significant cognitive and behavioral risk.

Attention Burden Score

81 / 100

Temu imposes one of the highest observed attention-burden profiles among e-commerce platforms analyzed under CAS methodology. The cumulative cognitive load is not incidental to the shopping experience; it is the product architecture. Users are simultaneously required to process time-depleting countdown timers, evaluate socially-proof-laden product listings, manage gamification reward states, navigate infinite content without structural endpoints, and resist persistent notification and re-enrollment nudges.

This multi-vector simultaneous stimulation overwhelms deliberative cognitive processing (System 2 reasoning) and systematically pushes users into fast, heuristic-driven (System 1) decision-making, which is precisely the neurological state most conducive to impulsive purchasing. Vulnerable users — including those with ADHD, anxiety disorders, financial stress, or limited digital literacy — face disproportionately higher effective burden due to reduced capacity to identify and resist these cumulative mechanisms.

Detected Patterns (10)

[HIGH] Infinite Scroll Feed

Infinite scroll removes natural stopping points from the browsing experience, eliminating the cognitive cue that typically signals session completion. This design decision is not neutral; it is engineered to suppress the user's ability to self-regulate session length by removing pagination boundaries.

[HIGH] Countdown Timer Pressure

Countdown timers are a classical scarcity-inducing dark pattern that artificially compress a user's decision-making window by creating time-bound urgency. Even when discounts are evergreen or automatically renewed, the visual presence of a depleting timer activates loss-aversion heuristics.

[CRITICAL] Variable Reward and Intermittent Reinforcement Loop

Variable reward schedules, in which rewards are delivered unpredictably, are among the most potent behavioral-conditioning mechanisms documented in psychological literature, comparable in mechanism to slot-machine reinforcement.

[HIGH] Attention Capture Mechanism

Pervasive attention-capture elements, including animated banners, flash-deal alerts, visual noise, badge notifications, and algorithmically prioritized high-CTR product placements, are deployed to override the user's intentional browsing agenda.

[MEDIUM] Social Proof Manipulation

Aggregated social-proof signals, including review counts, purchase-volume indicators, popularity badges, and 'others are viewing this' notifications, are deployed to manufacture perceived consensus and reduce individual deliberation.

[MEDIUM] Push Notification Coercion

Push-notification prompts, particularly when deployed at high frequency or strategically timed to interrupt user sessions, represent a form of re-engagement coercion. Combined with manipulative defaults that pre-select notification opt-in, they constitute a mechanism for sustained attention extraction.

[HIGH] Manipulative Default Settings

Manipulative defaults exploit status-quo bias by pre-selecting options that benefit the platform, such as marketing-communication opt-ins, data-sharing permissions, or promotional-subscription enrollment, requiring active user effort to reverse.

[HIGH] Roach Motel (Difficult Exit)

The roach-motel pattern describes an asymmetry between the ease of entering a service or commitment and the difficulty of exiting it. On e-commerce platforms this manifests as streamlined sign-up and subscription enrollment set against obfuscated cancellation paths and account-deletion complexity.

[HIGH] Forced Continuity

Forced continuity occurs when users are enrolled into recurring commitments, such as promotional memberships, auto-renewal subscriptions, or loyalty programs, without sufficiently prominent disclosure of the ongoing nature of the commitment, or when cancellation is made structurally difficult.

[CRITICAL] Gamified Engagement Loop

Temu's broader platform architecture incorporates game-like mechanics, including points accumulation, level progression, daily missions, referral challenges, and prize draws, that convert the shopping experience into a behavioral engagement loop structurally analogous to mobile gaming.

Recommendations

[IMMEDIATE] Disable or restructure countdown timers to reflect only genuine, non-resetting inventory or time constraints. Evergreen timers that reset upon expiry should be treated as deceptive urgency patterns and removed to align with emerging dark-pattern regulations in the EU (DSA), UK (Online Safety Act), and US FTC guidance.

[IMMEDIATE] Replace manipulative default opt-ins for push notifications and marketing communications with explicit, neutral opt-in prompts that present acceptance and refusal with equal visual prominence. Audit all checkout and account-creation flows for pre-selected consent fields.

[IMMEDIATE] Implement a clear, single-step cancellation and account-deletion pathway accessible from the user account dashboard without requiring customer-service contact or multi-step confirmation loops.

[SHORT TERM] Introduce session-boundary mechanisms, such as configurable time-on-site alerts, session-summary screens, or opt-in scroll-depth limits, to give users genuine agency over infinite-scroll session duration.

[SHORT TERM] Commission an independent audit of social-proof data sources and surface methodology, ensuring review counts and purchase-volume indicators reflect verified, locale-relevant data with clear provenance disclosure.

[SHORT TERM] Conduct a Child Safety Impact Assessment on gamification mechanics, variable-reward systems, and intermittent-reinforcement loops, and implement age-gate protocols for game-like reward features where legally required.

[LONG TERM] Restructure the core product-recommendation algorithm to optimize for purchase satisfaction and user-stated intent fulfillment rather than session duration and click-through-rate maximization.

[LONG TERM] Publish a transparent Attention Ethics Policy committing to measurable reductions in manipulative-pattern deployment, with annual third-party audits, public reporting, and proactive engagement with regulators under GDPR, DSA, and relevant APAC frameworks.

Analyzed: August 26, 2026, 6:47 PM EDT

9. DSA Risk Detection (EU Digital Services Act Module)

Run Date	Aug 26, 2026, 7:02 PM EDT
Methodology Version	DSA-RISK-V1.0
Input Mode	Crawl
Assessed Target	Temu — https://www.temu.com/
Run By	delante@asmcas.com
Section Status	Current (live DSA run for this assessment; 2 earlier runs retained for audit)

Total Systemic Risk

62.5 / 100 HIGH RISK

Total = IRS × 0.25 + MRS × 0.25 + UARS × 0.20 + SIRS × 0.20 + AS × 0.10. Classification bands: Low Risk 0–30, Moderate Risk 31–60, High Risk 61–80, Severe Systemic Risk 81–100. This run places the system in the High Risk band.

Systemic Risk Profile — Category Breakdown

Category	Score	Weight	Measures
IRS — Information Risk	61 / 100	25%	Virality rate (35%), amplification bias (40%), exposure skew (25%): how strongly the system distorts what information reaches a user.
MRS — Manipulation Risk	66 / 100	25%	Algorithmic weight (40%), engagement-loop strength (35%), recommendation dependency (25%): how much the system steers behavior rather than serving intent.
UARS — User Autonomy Risk	54 / 100	20%	Default bias (30%), control access (35%), opt-out barriers (35%): how far defaults and controls constrain a user's real choice.
SIRS — Societal Impact Risk	68 / 100	20%	Content clustering (30%), echo-chamber score (35%), exposure divergence (35%): population-level fragmentation and polarization pressure.
AS — Addictiveness	63 / 100	10%	Unweighted mean of scroll trigger, autoplay, infinite feed, and session duration: compulsion-forming design pressure.

Attention Harm Findings (Primary C.A.S. Outcome)

- Compulsion-forming design pressure (AS) — 63/100, elevated. Scored from scroll trigger 60, autoplay 51, infinite feed 92, and session duration 50. Higher values mean the interface itself sustains attention beyond user intent.

- Algorithmic steering of attention (MRS) — 66/100, elevated. Attention is allocated by the recommendation system rather than by the user, so exposure is a product of ranking, not choice.
- Loss of attentional autonomy (UARS) — 54/100, elevated. Defaults, control access, and opt-out friction determine how much attentional control the user actually retains.

DSA Compliance Flags: 0. Flags fire when a category reaches 70 or the total reaches 65; none were triggered by this run.

Input Risk Variables

Variable	Score
Virality rate	46.8 / 100
Amplification bias	65.1 / 100
Exposure skew	75.6 / 100
Algorithmic weight	60.7 / 100
Engagement loop strength	61.6 / 100
Recommendation dependency	80.5 / 100
Default bias	79.9 / 100
Control access	44.0 / 100
Opt-out barriers	43.0 / 100
Content clustering	78.5 / 100
Echo chamber score	58.1 / 100
Exposure divergence	69.6 / 100
Scroll trigger	60.0 / 100
Autoplay	51.0 / 100
Infinite feed	92.0 / 100
Session duration	50.0 / 100

Superseded DSA Runs (Audit History)

Run Date	Total Score	Classification	Mode	Superseded
Aug 26, 2026, 6:50 PM EDT	5.3 / 100	Low Risk	Crawl	Aug 26, 2026, 7:02 PM EDT
Aug 26, 2026, 6:48 PM EDT	5.3 / 100	Low Risk	Crawl	Aug 26, 2026, 6:50 PM EDT

Methodology: the DSA module scores 16 structural risk variables into five systemic risk categories (IRS, MRS, UARS, SIRS, AS) and a weighted total, using the linked assessment as its source context. DSA output is a governance-facing view of the same attention-harm evidence; it does not alter C.A.S., ASM, or SIA scoring, and it does not constitute legal advice.

Source: Assessment context — Temu, Temu Ecommerce, <https://www temu.com/>. Crawled source: <https://www temu.com/>. A full DSA report artifact reference is retained in the C.A.S. platform's linked file store for this assessment.

10. Evidence Log

Two evidentiary screenshots were uploaded in support of the AEA domain findings. Image files are retained in the source assessment record; this report references their metadata below.

File	Type	Domain	Uploaded	Notes
Screenshot, Aug 26 2026, 6:13:35 PM	Design	AEA	Aug 26, 2026	Documents the publicly visible Temu homepage at the time of assessment, showing the platform's visual hierarchy, promotional emphasis, engagement prompts, and pathways that encourage continued browsing or purchasing. Limited to visible website design; does not establish how the system operates for individual users.
Screenshot, Aug 26 2026, 6:11:08 PM	UX	AEA	Aug 26, 2026	Supporting UX capture for the AEA domain.

11. C.A.S. v2 Deterministic Scoring

Scored: August 26, 2026, 7:02 PM EDT

Domain	Score
AEA	68.3
AIS	52.0
SPI	45.1
CCS	85.0

Report Status & Composition Scope

Report State	Ready to finalize
Scoring Coverage	Not determined
Aggregate Score Status	Not available
Results Current	Yes

Scope: This report contains only current, non-archived artifacts linked to this assessment. Artifacts included: 0 requirement results, 0 statutory findings, 0 remediation tasks, 0 citation review items, 0 jurisdiction rollups.

12. Validity, Disclaimer, and References

Validity

This report reflects observed public-website conditions as assessed on and around August 26–27, 2026. Because the assessment status is Research Only, certification is false, and no certification date or validity end date has been issued, this document does not confer certified status and has no formal certification validity period. As a research record, it should be considered valid only for the assessed system state and evidence window, and it remains subject to revision, supersession by a later CAS assessment, material platform changes, or authority action.

Report Disclaimer

This CAS report is a research-based accessibility and autonomy assessment derived from the publicly observable behavior of the assessed system during the review window. It is not a legal ruling, security audit, consumer-protection enforcement action, or guarantee of future platform behavior. Scores and findings reflect the applied CAS methodology, sampled experience states, and available evidence at the time of assessment; they may change if the platform changes. The presence or absence of any signal, flag, or cohort impact in this report should not be used to make unsupported claims about protected groups beyond the limits of the documented method. This report should be interpreted as a structured compliance analysis intended to support remediation, governance, and further review.

References

Clark, D. (2026). Cognitive Attentional Standards (C.A.S.): An operational system for the evaluation of digital systems and measurement of attention harm.