

The Hyper-Personalization Readiness Worksheet

Hyper-personalization is personalization turned up, not a different thing. This is how to tell whether it will pay off or just cost you.

Step 1. Are the four requirements in place?

The technology rarely fails. The data, the volume, or the content usually does. Score honestly before spending.

Unified data

Our customer data is joined up across systems, current, and accurate. Usually a CDP.

1

2

3

4

5

Real-time capability

Our systems can process a signal and change the experience within one session.

1

2

3

4

5

Enough volume

We have the customer numbers for small per-person gains to multiply into real money.

1

2

3

4

5

Content and offer variety

We have enough content, offers, and range that meaningfully different experiences are possible.

1

2

3

4

5

Reading it: All four strong, hyper-personalization can pay off. Any one weak, fix that foundation first. Weak on data or content, do not start: aiming precisely and then showing everyone the same thing, on a fragmented picture, is worse than generic.

Step 2. Do you actually have a relevance problem?

The question is never 'should we hyper-personalize'. It is 'do we have a relevance problem at a scale that justifies this'.

The relevance problem we are trying to solve, in plain words

If you cannot name one beyond 'it sounds advanced', that is the answer.

Step 3. Which failure mode are you most exposed to?

Failure mode	Our exposure (H/M/L)
Crossing the creepy line	
Confident mistakes on imperfect data	
Stale signals treated as lasting preference	
Hidden ongoing data-engineering cost	

Two tests to keep running. The explain-it-out-loud test: would the person be comfortable if you told them why they are seeing this? And the comparison test: does this beat good ordinary personalization by enough to justify the cost, not just beat a generic experience? Often it does not, and finding that out cheaply is a win.