

A Source-Grounded Climate-Risk Outreach Protocol

Three customer-discovery messages for RelayProof

Spaces Cold Outreach Agent · Sender: Elena Park · 7 August 2026

Abstract

This document specifies a small, evidence-preserving cold-outreach campaign for RelayProof, an early concept for recording the evidence, ownership, and approval path behind climate-resilience investments. Three public leaders were selected because their organizations address the transition from climate-risk analysis to decision-ready action. For each prospect, the campaign separates the verified public hook from the sender’s hypothesis, uses one narrow 20-minute customer-discovery ask, and avoids guessed email addresses, invented referrals, attachments, or implied familiarity. The deliverable includes the prospect rationale, send-ready copy, follow-up rule, and pre-send verification protocol. Success is defined as learning about the handoff between risk intelligence and funded action, not booking a product demonstration or claiming endorsement.

1 Campaign Objective

RelayProof is testing a workflow hypothesis: after a credible climate-risk result exists, asset teams may still lack a durable record of assumptions, owners, approvals, expected avoided loss, and later evidence. The outreach seeks expert descriptions of where that handoff fails in practice.

The campaign has one call to action: a 20-minute customer-discovery conversation. It does not ask for a pilot, partnership, investment, referral, or product evaluation.

2 Selection and Evidence Rules

Prospects were selected on three criteria: public relevance to physical climate risk, visible attention to decision or investment workflows, and plausible ability to describe institutional failure modes. Every role and personalization hook must be traceable to an official organization page at the time of sending.

The message construction function can be represented as

$$M = H + G + Q, \quad (1)$$

where H is one verified public hook, G is the honest connection to the sender’s research gap, and Q is one proportionate question. A message fails validation if any term requires private knowledge or unsupported inference.

3 Prospect Set

No personal email address is inferred. Use the official organization contact route or a verified pro-

fessional profile. A public role is not permission to invent familiarity.

4 Message 1: Jupiter Intelligence

4.1 Subject: After the risk score

Rich,

Jupiter’s emphasis on translating climate exposure into defensible financial decisions gets at the part of climate software I find most consequential: what happens after the model is right.

I am building RelayProof, a lightweight evidence room for resilience decisions. We are testing whether asset teams need a clearer record of the assumptions, owners, approvals, and expected avoided loss behind each adaptation project.

I am not asking for a product pitch or partnership. Would you be open to a 20-minute conversation about where institutional clients still lose momentum between climate intelligence and funded action?

Best,

Elena Park

4.2 Audit note

The message connects an official positioning theme to one research question. It does not claim that Jupiter has a specific workflow failure or that the recipient has previously expressed interest in RelayProof.

Table 1: Prospect rationale and permitted personalization. Current role and source must be rechecked immediately before sending.

Prospect	Public relevance	Personalization hypothesis
Rich Sorkin, Jupiter Intelligence	Climate-risk intelligence and financial decision support	Ask what evidence remains missing after a high-quality risk model reaches a capital committee.
Lukky Ahmed, Climate X	Physical-risk analytics and decision-ready action for real estate and finance	Ask whether the handoff fails in analysis, organizational ownership, or budget approval.
Matthew Eby, First Street	Property-level climate-risk information and action-oriented risk communication	Ask where asset managers lose momentum between a location-level result and an operating or capital decision.

5 Message 2: Climate X

5.1 Subject: From analysis to approval

Lukky,

Climate X’s emphasis on decision-ready action rather than high-level indicators described the handoff I am researching almost exactly.

I am testing RelayProof, a small workflow that turns an existing climate-risk output into an approval record for one resilience project: evidence, cost assumptions, owner, decision, and later proof. Could I ask you 20 minutes of questions about where real-estate teams get stuck after they understand exposure? I am specifically trying to learn whether the gap is analytical, organizational, or simply budget ownership.

Best,

Elena Park

5.2 Audit note

The copy uses a public product theme without attributing a quotation that has not been preserved verbatim. The ask is for expertise, not endorsement.

6 Message 3: First Street

6.1 Subject: What blocks the action?

Matthew,

First Street’s work on moving from understanding physical climate risk toward action on real investment decisions stood out. The data layer is advancing quickly; the internal approval path appears far less standardized.

I am the founder of RelayProof, an early concept for preserving the evidence behind a resilience decision—what changed, who approved it, which assumptions mattered, and how the expected

benefit will be revisited.

Would you be open to a 20-minute customer-discovery call about where asset managers lose momentum between a location-level result and a capital or operating decision? I am looking for failure modes, not a sales opportunity.

Best,

Elena Park

6.2 Audit note

The message marks RelayProof as an early concept and the approval-path claim as the sender’s interpretation. It does not imply access to First Street customer data.

7 Follow-Up Protocol

Send one follow-up after five to seven business days. The follow-up should add one new observation or sharpen the research question. It should not say only “just bumping this.” If there is no response after the follow-up, close the sequence without another message.

A valid follow-up for this campaign is:

One detail I am trying to separate is whether adaptation decisions stall because the expected benefit is hard to express, or because no one owns the approval path. If that distinction shows up in your work, I would value 20 minutes to understand it. If not, no reply is needed.

8 Pre-Send Verification

Before sending, confirm the recipient’s current role, the exact public source supporting the hook, the

official or verified channel, and whether the organization has recently changed its positioning. Remove the message if the hook cannot be verified. Confirm that the copy contains no guessed address, invented referral, hidden tracking claim, attachment, calendar demand, or false urgency.

9 Measurement

Track delivery, reply, accepted conversation, and substantive learning separately. A reply is not automatically a successful outcome; the goal is evidence about the decision workflow. After each conversation, record the failure mode described, the role that owns the next step, the evidence required for approval, and what currently preserves the decision record.

10 Open Questions

The campaign has not yet established which asset class has the sharpest pain, whether the workflow owner sits in risk, finance, sustainability, or operations, or whether the missing artifact is analytical evidence, organizational accountability, or auditability. Those are the intended outputs of discovery.

11 Conclusion

The three messages are intentionally narrow. Each contains one verified hook, one honest research connection, and one modest ask. That restraint protects factual accuracy and makes the campaign useful even if no prospect becomes a customer.

Source Note

Public hooks were developed from official Jupiter Intelligence, Climate X, and First Street organization pages. Current roles, wording, and preferred channels must be rechecked immediately before sending. No private contact information is included in this example.