



NORTHHOLD
ADVISORY

JURISDICTION RISK MATRIX

Methodology and User Guide

VERSION 1.0

Effective 21 July 2026

255 jurisdictions covered

northholdadvisory.com

Contents

1. What this model is for - and what it is not
2. How the model decides: three tiers
3. The indicators, and why each one is here
4. Weighting: pricing in the overlap
5. Thresholds and calibration
6. Data sufficiency: when the model refuses to answer
7. Offshore centres
8. Override and rationale
9. Governance, review and change control
10. How models like this fail - and how to test yours
11. What a supervisor will ask
12. Adopting this in your firm

Annex A Source register

Annex B Jurisdiction identifiers and naming

Annex C Worked examples

How to read this guide

If you want to use the file, read sections 1, 2, 6, 8 and 12.

If you want to change the model, read 3, 4, 5 and 7.

If you are preparing for an audit or a supervisory visit, read 9, 10 and 11.

1. What this model is for - and what it is not

The Jurisdiction Risk Matrix classifies countries and territories for AML/CFT purposes into Low Risk, Medium Risk, High Risk, Blacklist, or Manual Review. It supports a risk-based approach: the rating is an input to customer and transaction risk assessment, not a decision in itself.

Where the output is used

- **Onboarding and KYC** - defining the depth of initial and ongoing due diligence.
- **Periodic review** - setting review frequency and documentation standards.
- **Transaction monitoring** - calibrating geographic rules and alert thresholds.
- **Product eligibility** - determining eligibility for particular products or corridors.
- **Third-party and vendor risk** - assessing counterparties with cross-border exposure.

Which jurisdictions to test

Assess every jurisdiction with a connection to the relationship, not only the country of incorporation: domicile and principal operations, payment origin and destination, the location of ultimate beneficial owners and controllers, intermediary and correspondent banks, and any place of physical presence, outsourcing or data storage. Where several jurisdictions are in scope, the highest applicable rating governs until due diligence resolves it.

What this model does not do

- It does not screen names. Sanctions and PEP screening remain separate mandatory controls; a country rating tells you nothing about the individual in front of you.
- It does not replace judgement. The model produces a defensible starting point. The Compliance function owns the final rating and must be able to depart from the model with reasons.
- It does not stay current by itself. Lists move - FATF plenaries roughly three times a year, EU and OFAC continuously. An un-refreshed matrix is not a neutral document; it is a control that has quietly stopped working.
- It is not legal advice. It is a risk-assessment tool built on public sources.

2. How the model decides: three tiers

Classification is strictly hierarchical. Once a condition is met at a higher tier, no lower-tier logic runs and the rating is final. This ordering matters: legal prohibitions must never be capable of being averaged away by favourable index scores.

Tier 1 - Blacklist (absolute prohibition)

Set where the jurisdiction appears on any of:

- FATF blacklist - High-Risk Jurisdictions Subject to a Call for Action
- UN Security Council sanctions
- The sanctions regime of your own home jurisdiction
- Your internal blacklist

Why the home-jurisdiction column ships empty

A template cannot assume whose law binds you. A Canadian money services business, a Lithuanian EMI and a UK payment institution are each absolutely bound by a different regime, and each treats the others as strong signals rather than prohibitions. Populate column E of Data · Lists with the regime that is legally binding on your firm - that is what makes Tier 1 correct for you rather than for someone else.

Tier 2 - Automatic High Risk

Where Tier 1 does not apply, High Risk is assigned without index scoring if the jurisdiction appears on: the FATF greylist (Jurisdictions under Increased Monitoring); the EU list of high-risk third countries; or EU, UK, US (OFAC) or Canadian territory-level sanctions.

What counts as a jurisdiction-level sanctions designation

Sanctions are overwhelmingly imposed on named persons and entities, not on territories. If every country hosting a designated individual were flagged, most of the world would carry a sanctions flag and the signal would carry no information.

The flag is therefore set only where measures restrict the territory as such: a comprehensive embargo, or broad sectoral measures against the economy. Purely targeted listings do not set it.

Write this definition into your policy. An unstated criterion is an unauditable one - the supervisor's question is not "is this country sanctioned?" but "on what basis did you decide it counts?"

Tier 3 - Weighted index model

Where neither override applies, five indicators are each converted to a risk category - Low = 1, Medium = 2, High = 3 - and combined as a weighted mean. Weights are renormalised over the indicators actually available, so a publisher that does not cover a jurisdiction reduces confidence without distorting the result.

The mean is then banded into equal thirds of the 1–3 scale:

Weighted mean	Rating
below $5/3$ (≈ 1.667)	Low Risk
$5/3$ up to below $7/3$	Medium Risk
$7/3$ (≈ 2.333) and above	High Risk

Why the bands are written as fractions

$5/3$ and $7/3$ are exact thirds. Written as rounded decimals - 1.67 and 2.33 - they create an ambiguous edge precisely where real results land: with three or six indicators, means fall exactly on 1.6667 and 2.3333. Whether such a jurisdiction rates Low or Medium then depends on how the threshold happened to be rounded when the policy was typed. Express the cuts as fractions and the ambiguity disappears.

3. The indicators, and why each one is here

Indicator	Publisher	What it measures	Direction
Basel AML Index	Basel Institute on Governance	ML/TF risk composite; incorporates FATF mutual-evaluation effectiveness data	0–10, higher = riskier
Financial Secrecy Index	Tax Justice Network	Banking secrecy, beneficial-ownership and corporate transparency	0–100, higher = riskier
Corruption Perceptions Index	Transparency International	Perceived public-sector corruption	0–100, higher = cleaner
Global Terrorism Index	Institute for Economics & Peace	Frequency and severity of terrorism	0–10, higher = riskier
WGI Political Stability	World Bank	Likelihood of instability or politically motivated violence	0–100, higher = more stable

On the Financial Secrecy Index: use the Secrecy Score

The FSI publishes two numbers. The headline FSI rank weights secrecy by the jurisdiction's share of global financial services, so large economies dominate it for reasons of scale rather than opacity. The Secrecy Score measures opacity alone. For jurisdiction risk you want the Secrecy Score; taking the rank instead would tell you that the largest markets are the most secretive, which is a statement about size, not about laundering risk.

Two indicators deliberately excluded

- **Economic Freedom Index** - it measures tax burden, regulatory efficiency and government intervention. A low-tax, lightly regulated jurisdiction scores as economically free - which the model would read as low risk, even though that profile is also the offshore secrecy profile. The indicator can invert the signal it is supposed to give. It is also published by an organisation with a stated policy position, which is uncomfortable ground on which to defend a risk rating.
- **Global Peace Index** - it already contains terrorism impact as a component, alongside conflict and militarisation. Running it beside the Global Terrorism Index double-weights conflict against financial-crime factors. The GTI is the more directly relevant of the two for terrorist financing, so it stays and the broader peace measure does not.

4. Weighting: pricing in the overlap

Indicator	Weight	Reasoning
Basel AML Index	40%	The only directly ML/TF-specific composite available, and the only one carrying FATF effectiveness data.
Financial Secrecy Index	20%	Secrecy is a sharper laundering driver than perceived corruption.
Corruption Perceptions Index	15%	Reduced - partially embedded within Basel.
Global Terrorism Index	15%	Least overlap with Basel; carries the terrorist-financing dimension.
WGI Political Stability	10%	Reduced - partially embedded within Basel.

The overlap problem, stated plainly

The Basel AML Index is a composite. Among its inputs are corruption measures, financial-secrecy measures and governance indicators - that is, the very indices sitting beside it in this model.

Give all five equal weight and corruption is counted twice, secrecy twice, political stability twice. The model looks balanced across five sources while leaning heavily in one direction.

The response here is not to hide the overlap but to price it: Basel carries the dominant weight, and its constituents are deliberately reduced. Document it as a known and managed limitation. A supervisor is not troubled by a limitation you have identified and addressed; a supervisor is troubled by one you appear not to have noticed.

All weights live on the Config sheet and are read by every formula. Setting all five to 20 reproduces a conventional unweighted model, if that is what your policy prefers. Nothing in the workbook hard-codes a weight.

5. Thresholds and calibration

Indicator	Low	Medium	High
Basel AML Index	≤ 3.99	≤ 5.99	> 5.99
Financial Secrecy Index	≤ 58	≤ 68	> 68
Corruption Perceptions Index	> 59	≤ 59	≤ 39
Global Terrorism Index	≤ 2	≤ 4	> 4
WGI Political Stability	≥ 76	≥ 60	< 60

Basel, CPI and GTI cuts follow established practice. The FSI and WGI cuts are set at the 33rd and 67th percentiles of the current published distribution, which places roughly a third of covered jurisdictions in each band and is straightforward to explain.

A live example of why calibration must be revisited

The World Bank's 2025 revision of the Worldwide Governance Indicators retired percentile ranks and replaced them with an absolute 0–100 scale anchored to fixed benchmark countries, recalculating history back to 1996.

The trap

Percentile ranks are uniform by construction, so a cut at the 16th percentile always separates the same proportion of countries. The replacement scale is absolute and therefore not uniformly distributed.

Carry old percentile thresholds across to the new scale and the model keeps running, keeps producing plausible ratings, and is quietly miscalibrated. Nothing breaks visibly. This is the characteristic failure of index-driven models: the numbers still arrive, so nobody looks.

Whenever a publisher revises methodology, recalibrate against the new distribution and record what you did in the Change Log. The Source Register carries a methodology-vintage column for exactly this reason - a date alone will not tell you the scale changed underneath you.

The corresponding upside: the absolute scale is better suited to risk work. Percentile ranks are relative, so a jurisdiction could only improve if others deteriorated. On an anchored scale, a year-on-year movement reflects an actual change in conditions.

6. Data sufficiency: when the model refuses to answer

No publisher covers every jurisdiction, and coverage is thinnest precisely where risk is often highest: small offshore centres and disputed territories. In the shipped dataset, all five indicators are available for 95 of 255 jurisdictions; 34 have none at all. How a model handles that gap determines whether it is a control or a decoration.

Three rules

1. **No indicator available** - the jurisdiction is classified Manual Review. It is never rated Low by default. Absence of data is not evidence of safety.
2. **Exactly one indicator available** - the resulting category is raised one level (Low becomes Medium, Medium becomes High). A single datapoint should not be allowed to understate risk.
3. **Coverage below the minimum** - if the indicators present account for less than half of total weight, the jurisdiction routes to Manual Review regardless of what they say.

Why coverage is measured by weight, not by count

Counting indicators treats them as interchangeable. They are not: under this model Basel alone carries 40% of the decision, while Political Stability and Terrorism together carry 25%.

A jurisdiction with those two lighter indicators and nothing else has two of five - which a counting rule would happily score - but only a quarter of the evidence the model was designed around. Measuring coverage by weight makes the confidence test match the model's own structure.

The practical consequence is that a substantial number of jurisdictions will land on Manual Review. That is the intended behaviour. A rating the evidence cannot support is worse than an honest referral to a human, because it looks like a control while providing none.

7. Offshore centres

An offshore designation sets a floor: the jurisdiction cannot rate below Medium. It does not override, and it is not added into the score.

The distinction matters. A common design treats offshore status as an additional High-risk input and averages it in with the indices. That was defensible when nothing else in the model measured opacity. With the Financial Secrecy Index present, secrecy is already measured directly and in detail - adding an offshore point on top counts the same characteristic twice, and does so invisibly, because the double-count is buried inside an average.

A floor achieves the intent - an offshore centre should not be waved through as Low Risk - without corrupting the arithmetic. The floor level is set on the Config sheet and can be raised to High if your risk appetite requires it.

8. Override and rationale

Compliance may override any automatic rating. The model informs the decision; it does not make it.

An override without a written rationale is an exception, not a decision. The workbook treats it as such: the Dashboard counts overrides missing a rationale, and that figure must be zero before the file is relied upon as evidence.

What a good rationale contains

What the model produced, and what you are changing it to.

The specific evidence relied on - a mutual evaluation finding, a supervisory action, adverse media, your own transaction experience with the corridor.

Why that evidence outweighs the model's inputs in this case.

Who approved it, and when it will be revisited.

"Local knowledge" and "commercial relationship" are not rationales. They are conclusions with the reasoning removed.

9. Governance, review and change control

Element	Standard
Ownership	The Compliance Officer owns the model and approves the methodology and any material change.
Event-driven refresh	Whenever a source changes: FATF plenary outcomes, EU delegated regulations, sanctions updates.
Scheduled review	Methodology reviewed at least annually; index thresholds re-checked when a publisher revises methodology.
Change control	Every change to data, thresholds or logic recorded in the Change Log with author, reason and approver.
Versioning	Version and effective date on the file; superseded versions retained for audit.
Independent testing	The model is periodically tested against its own written policy - see section 10.

Refresh cadence in practice

FATF - three plenaries a year, typically February, June and October. Lists can change at each.

EU high-risk third countries - follows FATF but on its own timetable and with a lag; the two lists are routinely out of step. Treat them as separate sources, never as one combined flag.

EU, UK, OFAC, Canada, UN sanctions - continuous. Quarterly review is a floor, not a target.

Indices - annual or biennial. Check the methodology vintage, not just the publication year.

10. How models like this fail - and how to test yours

Country-risk models rarely fail loudly. They keep producing ratings that look reasonable while the reasoning underneath has drifted. Below are the failure modes worth testing for, and a self-test for each. None takes long; all of them are things an examiner may find first.

Failure 1 - The spreadsheet and the policy disagree

The policy is written by people, in words. The workbook is built separately, in formulas. Over time the two drift apart, and because both look authoritative, nobody compares them. The usual culprit is rounding: a policy states thresholds as decimal cut-points while the spreadsheet uses a rounding function, which is a different model with different boundaries.

Self-test - ten minutes

1. Take your policy's stated aggregation rule and implement it independently, in a second column or a short script.
2. Run both across every jurisdiction and compare the outputs, not the inputs.
3. **Count the disagreements and look at where they cluster.**

If any jurisdiction is rated differently by your tool than by the policy that governs it, you have an undocumented model. Fix the tool, or amend the policy - but do not leave both in force. Pay particular attention to disagreements in the conservative direction: a tool rating a jurisdiction lower than your policy requires is the finding that matters.

Failure 2 - Country names that silently mismatch

Every source names countries differently, and two pairs cause real damage: the Democratic People's Republic of Korea against the Republic of Korea, and the Democratic Republic of the Congo against the Republic of the Congo. Text matching that strips words like "Democratic" and "Republic" as noise will collapse each pair - attaching one country's designation to another's risk profile, with no error message.

Self-test

Join every source on ISO 3166 alpha-3 codes, never on names. Where a source supplies only names, map them once, deliberately, and keep the mapping under version control.

Then check three things: that both Koreas and both Congos carry distinct codes and distinct ratings; that British and US Virgin Islands are separated; and that your row count matches the source's published count exactly. A single missing row is usually a name that failed to match and was dropped in silence.

Failure 3 - Hidden columns

Input columns get hidden to make a sheet readable, and the hidden logic stops being reviewed. It is the mechanism by which a model's assumptions become invisible to the people accountable for them. Separate inputs onto their own sheet instead. Nothing in this workbook is hidden.

Failure 4 - A stale list nobody noticed

Jurisdictions leave the FATF greylist as well as join it. A matrix refreshed on joiners but not on leavers over-rates countries for years, quietly generating enhanced due diligence - and cost, and customer friction - that your own policy does not require.

Self-test

For each list, compare your flags against the publisher's current list in both directions: who have they added that you lack, and who have they removed that you still carry?

Record an as-of date per source, not one date for the whole file. "Last updated" on a workbook tells a supervisor nothing about which edition of which list you actually applied.

11. What a supervisor will ask

The question	What a good answer looks like
Why does this jurisdiction have this rating?	Name the rule that fired. The Matrix states it per row - an override tier, or the weighted model with its inputs.
Which edition of the FATF list did you apply?	The Source Register carries an as-of date and edition for every source, not a single file date.
Why do these indicators, and not others?	Section 3 - including the two deliberately excluded and the reasons.
Why is this one weighted more heavily?	Section 4 - relevance to ML/TF, and the documented overlap adjustment.
Are you double-counting anything?	Yes, partially, and it is priced in. Naming the limitation is the answer; denying it is not.
What happens when data is missing?	Manual Review, never a default Low. Section 6.
Who can override, and where is the reasoning?	Compliance; rationale mandatory; the Dashboard reports any override lacking one.
When did you last review this, and who approved it?	The Change Log.
How do you know the tool matches the policy?	You tested it, and the test is recorded. Section 10.
What are the model's weaknesses?	Answer directly and specifically. An examiner trusts a model whose limits you can state faster than they can.

12. Adopting this in your firm

1. Decide whose sanctions bind you absolutely and populate the home-jurisdiction column. Until you do, Tier 1 reflects only FATF and the UN.
2. Review the weights on Config against your own risk appetite and business model. A firm concentrated in crypto corridors may justify a heavier terrorist-financing weighting; write down whatever you choose and why.
3. Confirm the thresholds, and recalibrate any index whose publisher has changed methodology since this edition.
4. Populate your internal blacklist.
5. Refresh every source from the primary references in Annex A and record the as-of dates.
6. Work through the Manual Review queue. Document each decision as you would an override.
7. Run the self-tests in section 10 and keep the results. That record is the difference between asserting the model works and demonstrating it.
8. Approve the model, record it in the Change Log, and set the next review date.

A closing word on proportionality

This model is deliberately conservative where evidence is thin, because in AML the cost of a false negative is asymmetric. But conservatism has a price too: every unnecessary High rating consumes analyst time and customer goodwill. If your matrix rates most of the world High, it is not cautious - it is uninformative, and your team will start working around it.

The objective is a model that discriminates: that says Low with confidence where the evidence supports it, says High where it matters, and admits when it does not know.

Annex A Source register

Source	Publisher	Role in model	Cycle
FATF black & grey lists	FATF	Tier 1 / Tier 2	~3 plenaries a year
EU high-risk third countries	European Commission	Tier 2	Ad hoc; lags FATF
UN Security Council sanctions	United Nations	Tier 1	Continuous
Home-jurisdiction sanctions	Your regulator	Tier 1	Continuous
EU / UK / OFAC / Canada sanctions	Respective authorities	Tier 2	Continuous
Offshore financial centres	Eurostat	Floor rule	Infrequent
EU non-cooperative tax jurisdictions	Council of the EU	Context	Twice a year
Basel AML Index	Basel Institute on Governance	Index - 40%	Annual
Financial Secrecy Index (Secrecy Score)	Tax Justice Network	Index - 20%	Biennial
Corruption Perceptions Index	Transparency International	Index - 15%	Annual
Global Terrorism Index	Institute for Economics & Peace	Index - 15%	Annual
WGI Political Stability	World Bank	Index - 10%	Annual

Full URLs, editions and as-of dates are maintained on the Source Register sheet of the workbook, which is the controlled record. This annex is a summary.

Annex B Jurisdiction identifiers and naming

The matrix is keyed on ISO 3166-1 alpha-3 codes, with alpha-2 and numeric codes alongside so it joins cleanly to onboarding and screening systems. Names are for humans; codes are for machines.

Six territories carry no ISO 3166-1 code but must still be assessed: Kosovo, Abkhazia, South Ossetia, Transnistria, Northern Cyprus and Nagorno-Karabakh. They are assigned codes from the user-assignable range that ISO reserves for exactly this purpose, and are marked "User-assigned (non-ISO territory)" in the ISO status column so nobody mistakes them for official codes. Kosovo uses XK / XKX, the convention already followed by the European Commission, IMF and World Bank.

Note that some publishers use their own country codes which do not match ISO - the World Bank, for instance, codes Andorra as ADO where ISO uses AND. Map on names once, deliberately, then key on ISO thereafter.

Annex C Worked examples

United Kingdom - how a Low Risk rating is built

Indicator	Value	Category	Weight	Contribution
Basel AML Index	4.04	Medium (2)	40%	0.80

Financial Secrecy Index	45	Low (1)	20%	0.20
Corruption Perceptions Index	70	Low (1)	15%	0.15
Global Terrorism Index	2.94	Medium (2)	15%	0.30
WGI Political Stability	70.3	Medium (2)	10%	0.20
Weighted mean			100%	1.65

1.65 falls below 5/3 (1.667), so the rating is Low Risk. It is worth noticing how narrow that margin is - and that a model using a rounding function rather than stated bands would place the same jurisdiction in Medium Risk. Two defensible-looking designs, different answers, thousands of relationships affected.

Other patterns

- **Iran** - FATF blacklist fires at Tier 1. No scoring is performed, and favourable index values are irrelevant - which is the point of the hierarchy.
- **British Virgin Islands** - only one indicator is available, but the jurisdiction appears on both the FATF greylist and the EU high-risk list. Tier 2 resolves it as High Risk before data sufficiency is ever tested.
- **Cayman Islands** - two indicators present, together carrying 30% of total weight - below the 50% minimum. Routed to Manual Review rather than scored on partial evidence.
- **Philippines** - all five indicators present, weighted mean 2.40, above 7/3. High Risk. Under a rounding-based aggregation the same inputs produce Medium - an under-rating, and the direction of error that matters most.

Northhold Advisory · Senior-led AML & financial-crime advisory · EU · UK · CIS

northholdadvisory.com

This guide and the accompanying workbook are provided for general information and do not constitute legal, regulatory or financial advice. Data is drawn from public sources as at the dates recorded in the Source Register and must be verified before operational use. Northhold Advisory provides compliance advisory services only and accepts no liability for reliance on these materials.