

The Glass Never Lies

A Physics Framework for Reading and Projecting Whisky from First Principles

WhiskyBaba Research

February 2026

Executive Summary

Every bottle of whisky contains a physical record of its own history. The chemistry is not ambiguous — oak extraction follows documented kinetics, cask type determines specific phenolic profiles, and colour develops along predictable curves governed by temperature, wood species, fill number, and time. The science has been published for decades. What hasn't existed is a framework that reads this record systematically.

This paper introduces a bidirectional physics framework for whisky liquid intelligence. The concept is straightforward: if production creates chemistry in documented ways, that chemistry can be interrogated in both directions.

In the **inverse direction**, a photograph of whisky liquid yields colour data. That colour encodes information — hue indicates cask influence, saturation reflects extraction intensity, and lightness constrains age range within a given cask regime. The approach doesn't replace laboratory analysis. It provides a non-destructive screening layer that currently doesn't exist between expert opinion and mass spectrometry.

In the **forward direction**, production parameters — cask type, oak species, fill number, warehouse climate, target ABV — feed into the same underlying physics to project maturation trajectory. A distiller commits wood and time measured in decades. The chemistry that governs colour development also governs flavour compound accumulation. Projecting outcomes from documented physics reduces the blind gamble without eliminating the craft.

The framework rests on peer-reviewed science: oak extraction kinetics in wine and spirits fit Peleg-type models, with parameters published for several relevant cask and solution conditions; European and American oak produce measurably different phenolic fingerprints; colour correlates strongly with chemical maturity within controlled systems. The framework draws on published work in three areas — extraction kinetics, cask-type differentiation, and fill-number effects — and adds the systematic exploitation of these relationships, along with the honest acknowledgment of where they break down.

Colour alone cannot determine age. The same amber can emerge from six years in first-fill Pedro Ximénez sherry or eighteen years in refill bourbon. Without cask information, uncertainty ranges are wide. With cask type and fill number known, they tighten considerably. The system's value scales with available metadata, and it says so.

The inverse and forward directions share the same underlying physics model. Success in one direction validates the other — a structural property that pure statistical approaches cannot replicate. This bidirectional validation is the framework's distinctive contribution.

Two proof-of-concept systems demonstrate feasibility. Neither is presented as finished product. Both are presented as evidence that the physics is exploitable. What comes next — systematic calibration with known-parameter samples, tighter confidence ranges, independent validation — is the work this paper exists to enable.

1. The Problem: Whisky Lacks Liquid-Level Intelligence

Pour a glass of whisky. Hold it to the light. The liquid tells a story — amber from oak tannins, copper hints from sherry influence, pale gold from long years in exhausted wood. Every serious taster reads these cues instinctively. None of them reads them systematically.

The whisky industry has built sophisticated infrastructure around everything except the liquid at the point where decisions are made. Provenance documentation tracks ownership chains. Packaging authentication catches counterfeit labels. Expert panels evaluate sensory character. Auction houses verify seals and fill levels. Distillery laboratories already run UV-vis, GC-MS, and NIR on production samples; the chemistry is watched closely inside the gates. But in the auction room, the warehouse, the retailer's shelf — the places where purchase and triage decisions actually happen — there is no physics-based interrogation of the liquid at all.

This is not because the science doesn't exist. Oak extraction chemistry has been documented for decades. Ellagitannin kinetics are published. Cask type colour signatures are measurable. The relationship between production parameters and chemical outcomes is well-characterised in peer-reviewed literature. The gap isn't knowledge. It's application.

The cost of this gap is quantifiable. Cook et al. (2020) subjected 55 bottles of "rare" Scotch whisky to radiocarbon dating at the Scottish Universities Environmental Research Centre. Twenty-one failed — either entirely modern spirit sold as pre-1900, or significantly misdated. That's a 38% failure rate among bottles commanding premium prices, tested by arguably the most rigorous method available (DOI: 10.1017/RDC.2019.153). The Rare Whisky 101 index, tracking the secondary market, estimates £41 million in suspect bottles circulating at any given time.

Radiocarbon dating works. It also requires physical sampling, specialist laboratory access, weeks of turnaround, and costs that make it impractical for routine screening. Mass spectrometry identifies chemical fingerprints with precision. It also destroys the sample. The whisky industry's authentication toolkit is powerful at the extremes — definitive when deployed — but entirely absent at the everyday level where most purchasing decisions happen.

Expert tasting, the field's default assessment method, carries its own documented limitations. Morrot, Brochet, and Dubourdieu (2001) demonstrated that trained oenology students, given identical white wines with one sample dyed red, described the "red" wine using red-fruit vocabulary and the "white" using citrus-floral terms. The same liquid, different visual cues, completely different expert descriptions (DOI: 10.1006/brln.2001.2495). Quigley-McBride et al. (2018) found that whisky tasters' preferences shifted based on presentation order — serial position effects overriding intrinsic quality assessment (DOI: 10.3758/s13420-018-0346-9). Gawel and Godden (2008) reported that only 32% of experienced wine judges produced stable individual rankings across repeated sessions (DOI: 10.1111/j.1755-0238.2008.00001.x).

This is not an indictment of expertise. Trained palates detect nuance that instruments miss. The point is different: human sensory assessment is demonstrably susceptible to cognitive bias from visual cues, presentation context, and fatigue. A physics-based screening layer would

not replace expert judgment. It would provide an objective baseline that doesn't shift with the order of service.

The most recent broad review of whisky authentication methods — Okolo et al. (2023), covering analytical advances from 2018 to 2023 — catalogues spectroscopic, chromatographic, and chemometric approaches, all of which require laboratory equipment, trained operators, and physical samples (DOI: 10.1111/1541-4337.13249). Chen et al. (2025) survey spirit authentication more broadly, noting the field's trajectory toward portable sensors integrated with artificial intelligence (DOI: 10.1016/j.foodchem.2025.145898). The direction is clear: the industry needs faster, less destructive, more accessible assessment tools. What neither review considers is the possibility that visible colour — the most immediately available physical signal — already encodes extractable intelligence when interpreted through documented physics.

That's the gap this paper addresses. Not the absence of analytical chemistry. The absence of a framework that treats the liquid's visible properties as a physics problem with two exploitable directions.

2. The Science: What Colour Carries

New make spirit is colourless. Every shade of gold, amber, copper, and mahogany in the glass arrived there through interaction with oak during maturation. This is not metaphor. It is extractive chemistry with documented kinetics, identifiable compounds, and measurable relationships to production parameters.

Understanding what colour encodes requires following the chemistry through three domains: which compounds create colour, how those compounds accumulate over time, and why different production regimes produce different colour outcomes.

2.1 The Compounds

Whisky colour derives primarily from four compound classes extracted from oak wood, each contributing distinct chromophoric properties.

Ellagitannins — castalagin, vescalagin, and their derivatives — are the dominant initial colour contributors. These are large polyphenolic molecules (molecular weight 900–1,100 Da for monomers) that dissolve rapidly from the wood surface during early maturation. Their behaviour over time is counterintuitive and important: concentration rises steeply in the first one to three years, then *declines* as hydrolysis and oxidation convert them to smaller fragments and brown polymeric structures (Viriot et al., 1993; DOI: 10.1021/jf00035a013). The colour deepens even as ellagitannin concentration falls — because the oxidation products are more intensely coloured than the parent compounds.

Phenolic acids — gallic acid, ellagic acid, and related structures — accumulate more steadily. These lower-molecular-weight compounds (170–302 Da) contribute yellow-brown hues and continue to increase throughout maturation without the peak-and-decline pattern of ellagitannins (Fujieda et al., 2008; DOI: 10.1021/jf8012713).

Furanic aldehydes — furfural and 5-hydroxymethylfurfural — form from carbohydrate degradation in toasted or charred wood. Their concentration correlates with toast level and contributes caramel and golden-brown coloration. Heavily charred bourbon casks release measurably more furfural than lightly toasted European oak (Nocera et al., 2023; DOI: 10.1021/acs.jafc.3c00501).

Lignin-derived aldehydes — vanillin, syringaldehyde, coniferaldehyde — contribute brown colour development alongside their more famous role as aroma compounds. These accumulate throughout maturation, and their concentrations at any given age depend on oak species, toast level, and fill history (DOI: 10.3390/foods12173190).

A spectrophotometric study of 68 whiskies confirmed strong correlation between total phenolic content and colour density ($R^2 > 0.9$ within controlled cask regimes), indicating that colour intensity is a useful proxy for extractive load when the cask type and treatment are known (DOI: 10.1002/jsfa.6960).

2.2 The Kinetics

Colour doesn't accumulate linearly. A whisky doesn't gain the same increment of colour in year fifteen that it gained in year one. The extraction follows documented kinetic models — and the shape of the curve matters for any system attempting to infer age from colour.

Research identifies three distinct phases. In the first six months, rapid surface extraction dominates: readily soluble phenolics dissolve from the inner wood surface, producing a steep rise in both phenolic content and colour intensity. From roughly six months to three years, diffusion-controlled extraction takes over as compounds migrate from deeper wood layers — slower, but still substantial. Beyond three years, chemical transformation dominates: extraction rates decrease, but oxidation and polymerisation reactions continue converting extracted compounds into increasingly coloured products (Delgado-González et al., 2024; DOI: 10.3390/foods13203279).

The practical consequence is that simple linear models — “twice the age, twice the colour” — are inadequate. In wine and spirits studies, the data fit Peleg-type pseudo-second-order kinetics: rapid early accumulation approaching an asymptotic maximum, with rate constants that vary by cask regime and extraction conditions (DOI: 10.1016/j.foodcont.2021.107934). This non-linearity is what makes colour informative. A pale whisky in a first-fill sherry cask tells you something very different from the same paleness in a refill bourbon cask. The kinetic curves differ, so the age implications diverge.

2.3 Why Cask Type Changes Everything

Different cask regimes produce measurably different colour outcomes — not just in intensity but in *hue*. This is the chemical basis for cask-type inference from colour.

Bourbon casks (American oak, *Quercus alba*) contribute colour primarily through oak-native compounds: ellagitannins, furanics from the charred inner surface, and lignin-derived aldehydes. No wine-derived compounds are present. The resulting colour profile runs from golden to amber. In controlled conditions, bourbon-cask whiskies tend to cluster in higher hue ranges with relatively neutral a^* values in CIE $L^*a^*b^*$ colour space.

Sherry casks introduce an additional chromophoric layer. Oloroso-seasoned European oak (*Quercus robur*) extracts faster and at higher phenolic concentrations than American oak, and the wine residue contributes polyphenol glycosides — quercetin-glucuronide, myricetin-glucoside — that shift the colour toward red-brown hues (Roullier-Gall et al., 2018; DOI: 10.3389/fchem.2018.00029). Pedro Ximénez casks, with their higher residual sugar content, add Maillard reaction products that push colour further toward deep amber and mahogany.

Valcárcel-Muñoz et al. (2021) directly compared brandy aged in new oak versus sherry-seasoned casks, documenting measurably different total phenolic index values and colour density at equivalent aging times (DOI: 10.3390/foods10020264). Roullier-Gall et al. (2018) used ultra-high-resolution mass spectrometry on 150 whiskies ranging from one day to 43 years old, confirming that bourbon and sherry cask whiskies carry distinct molecular fingerprints detectable

through their different chemical profiles.

The “red channel” provides a practical differentiator. Higher red contribution (lower hue, positive a^* in CIE L*a*b* colour space) indicates wine or sherry influence. Lower red (higher hue, neutral a^*) indicates bourbon or plain oak maturation. This follows directly from the way wine-derived anthocyanin derivatives shift the spectrum when they’re present in one cask type and absent in the other. In practice, glass colour, dilution, chill filtration, and finishing regimes all add noise, so hue is a probabilistic signal rather than a categorical one.

2.4 Fill Number: The Depleting Resource

Each time a cask is filled, it surrenders extractable compounds. The relationship is steep and well-documented.

Puech (1990) quantified ellagitannin depletion across successive fills, showing 70–80% of extractable ellagitannins removed during first use (PMID: 2211469). Chira and Teissedre (2015) confirmed that refill wood extracts 40–60% less total phenolics compared to new or first-fill casks (DOI: 10.1016/j.foodchem.2013.01.039). Canas et al. (2022) tracked the progression further: vescalagin concentration dropped from approximately 15 mg/L in first extraction to below 4 mg/L by the third (DOI: 10.1016/j.lwt.2022.113664).

Across studies on spirits and model systems, second-fill casks yield approximately 30–50% of first-fill phenolic extraction; third-fill drops to 20–30%; by fourth use, extraction capacity is substantially diminished (Valcárcel-Muñoz et al., 2021; DOI: 10.3390/foods10071594). The exact percentages depend on species, seasoning history, and use conditions — these are approximate ranges, not universal constants. The colour implications are direct: the same spirit aged for the same duration will appear measurably lighter in a second-fill cask than a first-fill, and lighter still in subsequent fills.

This depletion curve is exploitable in both directions. In the inverse direction, unusually light colour for a stated cask type suggests higher fill number or shorter maturation. In the forward direction, known fill number allows projection of colour trajectory at any given age.

How Much of This Is Documented?

The framework does not invent new physics. It assembles published relationships that already exist in the literature.

Phenomenon	Quantitative relationship	Representative source
Colour vs time	Peleg-type and logarithmic models for colour and lightness evolution in spirits, with R^2 values typically above 0.9 within a given cask regime	Delgado-González et al. (2021), Food Chem. 360, 130001. DOI: 10.1016/j.foodchem.2021.130001
Cask size	Extraction rate scales with surface-to-volume ratio; quarter casks extract roughly 1.5–2.0 times faster than hogsheads for the same spirit	Trillo-Ollero et al. (2024), Food Chem. 140128. DOI: 10.1016/j.foodchem.2024.140128
Temperature	A 10 °C increase in maturation temperature increases extraction rates by roughly a factor of 1.8–2.0, consistent with Arrhenius behaviour	Jiménez-Moreno et al. (2019), Food Res. Int. 125, 108577. DOI: 10.1016/j.foodres.2019.108577
Colour vs phenolics	Absorbance at 420 nm and CIE b^* correlate strongly with total phenolic index (typical R^2 0.8–0.9)	Collins et al. (2016), J. Sci. Food Agric. 96, 4557–4567. DOI: 10.1002/jsfa.6960; Nocera et al. (2020), OENO One 54(3). DOI: 10.20870/oeno-one.2020.54.3.3114
Cask type fingerprints	Bourbon and sherry cask whiskies occupy distinct regions in a^*/b^* space and show different hue ratios and molecular fingerprints	MacKenzie & Aylott (2004), Analyst 129, 607–612. DOI: 10.1039/b403068k; Roullier-Gall et al. (2018), Front. Chem. 6, 29. DOI: 10.3389/fchem.2018.00029

These are the kinds of relationships the framework relies on: documented kinetics, measured colour–chemistry correlations, and cask-specific signatures rather than black-box correlations.¹

¹A fuller compilation of laws, constants, and citations sits behind the concept paper as internal research material.

3. The Inverse Direction: Inferring History from Colour

A glass of whisky sits on a table. No label visible. What can the liquid itself tell you?

Current practice offers three approaches. An expert can look at it, swirl it, nose it, and offer an opinion — informed by experience, subject to the biases documented in Section 1. A laboratory can take a sample, run chromatography or spectroscopy, and return precise chemical composition — destroying the sample, requiring specialist equipment, and producing results days or weeks later. Or the owner can trust whatever documentation accompanies the bottle.

The inverse direction proposes a fourth: photograph the liquid and apply documented physics to extract what the colour encodes.

3.1 What a Photograph Contains

A properly captured photograph of the liquid yields three measurable properties in HSL colour space: hue (the dominant wavelength — gold versus amber versus copper), saturation (colour intensity — vivid versus washed-out), and lightness (how dark or bright the liquid appears). Those values aren't subjective impressions; they're numerical coordinates that can be extracted, standardised for lighting conditions, and compared against documented chemistry.

Each coordinate carries its own kind of information. Hue indicates the chromophoric profile — bourbon cask maturation produces different dominant wavelengths than sherry cask maturation, because different compound classes are present. Saturation reflects extraction intensity — first-fill casks produce higher chromatic saturation than depleted refills at equivalent ages. Lightness constrains the age-within-cask-regime — darker means more extraction, which correlates with longer maturation, though the relationship follows the non-linear kinetics described in Section 2, not a simple proportional rule.

3.2 The Inference Chain

With colour coordinates extracted, the physics documented in Section 2 runs in reverse. The hue narrows the probable cask type. A hue in the 35–50° range, with minimal red contribution, is consistent with bourbon cask maturation. A shift toward lower hue values with increased red component points to sherry or wine cask influence. The saturation, combined with the hue-implied cask type, constrains the probable fill number — high saturation in a bourbon-regime hue suggests first-fill; the same hue at lower saturation suggests refill. Lightness, interpreted through the appropriate Peleg kinetic curve for the inferred cask regime, constrains the probable age range.

Each step narrows the probability space. Each step also carries uncertainty that must be reported honestly. The system outputs ranges, not point estimates, and confidence levels that reflect the genuine ambiguity at each inference stage.

3.3 Orphan Casks and Old Bottles

The place where this inverse direction bites hardest is not in everyday bottlings with clean paperwork. It is in orphaned stock and old bottles where records are thin, inconsistent, or plainly wrong.

For orphan casks, the combination of colour coordinates and basic physical parameters — declared region, measured ABV, known warehouse type — constrains what is chemically plausible. A five-year Indian single malt matured at 30 °C does not land on the same Peleg curve as a five-year Speyside malt matured at 12 °C; climate acceleration factors in the maturation literature suggest effective extraction rates that are roughly three times higher in the Indian case than in the Scottish reference.² Colour that looks “fifteen Scottish years” on the glass but sits in a cask declared as five-year Indian stock is not an anomaly — it is exactly what the physics predicts.

For historic bottles, the same logic runs through a different filter. The framework cannot certify authenticity, but it can compare what the liquid shows to what the label claims in three ways: whether the colour profile matches the claimed cask regime, whether the implied maturation rate is consistent with the declared region and era, and whether the colour behaviour under dilution is consistent with purely oak-derived colour rather than heavy caramel standardisation. Bottles that fail those physics checks are not automatically fake, but they move from “assume fine” into “worth sending to a lab”.

3.4 Confounders and Boundary Conditions

The inference chain described above works best under defined conditions: single cask regime, no E150a addition, no chill filtration that strips phenolics, and a photograph taken through clear glass under reasonable lighting. In practice, several confounders degrade the signal.

E150a spirit caramel adds colour without corresponding phenolic load — the system can flag likely non-oak colour through dilution behaviour (see below) but cannot subtract it from a static measurement. Blended or vatted whiskies combine multiple cask regimes into a composite colour that the framework isn’t designed to deconvolve. Chill filtration removes some phenolic material, potentially reducing colour intensity relative to what the cask regime would predict. Bottle age and storage conditions can shift colour over decades through slow oxidation.

The framework handles these by bounding its claims. When metadata is absent and confounders cannot be excluded, confidence ranges widen and the system says so explicitly. Useful inference is promised only within acknowledged bounds — single-regime or known-regime liquid, with stated assumptions about filtration and colouring practice.

²See, for example, Trillo-Ollero et al. (2024), Jiménez-Moreno et al. (2019), and regional angel’s share data summarised in the maturation literature.

3.5 The Dilution Behaviour Principle

One question that colour analysis confronts is whether the colour is entirely oak-derived or whether spirit caramel (E150a) — a legally permitted additive in many whisky-producing regions — has been introduced. E150a and natural oak colour overlap substantially in visible spectrum, making static colour measurement insufficient for distinction.

However, the physical chemistry of the two colour systems differs in ways that are potentially exploitable. E150a is a polydisperse melanoidin-type material formulated for solution stability across alcohol concentrations — it behaves as a simple water-soluble colorant under dilution (Licht et al., 1992; PMID: 1644378). Oak-derived phenolics, by contrast, participate in colloidal interactions sensitive to ethanol strength — they bind to tannin-protein complexes that respond non-linearly to changes in the solvent matrix (Fujieda et al., 2008; DOI: 10.1021/jf8012713; Mosedale & Puech, 1998; DOI: 10.1016/S0924-2244(98)00024-7).

On paper, that difference should translate into distinct dilution curves. The principle is established physical chemistry applied in a new context; the specific implementation and threshold calibration are part of ongoing validation work. No published study has yet demonstrated E150a versus oak colour separation via dilution alone in whisky-strength matrices — this remains a physically grounded hypothesis, not a validated detection method.

To be clear about what this does not do: it does not definitively detect E150a. It flags samples whose dilution behaviour is inconsistent with purely natural colour. The distinction between “flagging for further investigation” and “confirming adulteration” is important, and the system maintains it.

3.6 Positioning Against Existing Methods

The inverse direction does not compete with laboratory analysis. It occupies a different niche entirely.

Method	Requires	Destroys Sample	Turnaround
Expert opinion	Trained palate	No	Immediate
Colour-physics inference	Photograph	No	Seconds
Spectroscopy	Lab equipment	Partial	Hours–days
Radiocarbon dating	SUERC-class lab	Yes	Weeks

The value is not precision — laboratory methods will always be more precise. The value is accessibility, speed, and non-destructiveness at a point in the assessment pipeline where nothing currently exists. A collector considering a purchase at auction has no practical screening tool between “look at it and guess” and “send it to a lab.” This framework offers something in that gap.

4. The Forward Direction: Projecting Outcome from Production

If the inverse direction reads history from colour, the forward direction projects future from decisions. A distiller commits cask wood, warehouse space, and calendar time measured in decades. The chemistry that governs what those decisions produce is documented. The practical question is whether those documented relationships can be systematised into projection that's useful before the cask is opened.

4.1 The Production Parameter Space

Every whisky production decision has documented chemical consequences. Barley variety affects amino acid composition in the wash. Fermentation duration determines ester and higher alcohol profiles. Still shape and condenser type govern copper contact, which selectively removes sulphur compounds. Peat exposure introduces phenolic compounds at measurable concentrations. These pre-maturation decisions establish the new make character — the chemical foundation on which cask maturation builds.

Maturation introduces its own parameter set: oak species, cask type and previous contents, fill number, cask size (surface-to-volume ratio), warehouse environment, and intended duration. Each variable has peer-reviewed documentation of its chemical effect. European oak extracts faster and at higher phenolic concentrations than American oak. First-fill casks contribute measurably more than refills. Warmer warehouses accelerate extraction through well-established thermodynamic principles — reaction rates increase predictably with temperature, consistent with Arrhenius kinetics. In practice, distilleries don't control temperature with laboratory precision; they inherit whatever their climate and warehouse design give them. The model treats temperature as an average driver, not a precise knob. Smaller casks extract faster due to higher surface-to-volume ratios.

The forward projection system takes these parameters as input and calculates the accumulated chemical effect at any specified age, using the kinetic models and modifier relationships described in Section 2 and established in the peer-reviewed literature. At this stage, rate constants and modifiers come from published extraction studies, adjusted against a small internal set of known-parameter casks. The point of this concept paper is to show that even this first-pass calibration produces ranges consistent with real liquids; the heavy lifting on systematic calibration remains ahead.

4.2 Anchoring Projection in New Make

Forward projection is not just about cask choice and warehouse climate. It starts with the chemistry of the new make itself. Different distilleries produce systematically different new make: heavier cut points, longer fermentations, higher phenolic loads, and different ester balances all show up as measurable differences before the liquid ever sees oak.

The literature does not yet provide a closed-form mapping from new-make composition to later flavour, but it does show consistent patterns. Heavier new make with higher initial phe-

nolic content darkens faster and reaches higher total phenolic index for a given cask regime than lighter new make; longer fermentations and fruitier ester profiles change how oxygen and oak-derived aldehydes are expressed in the glass.³ The forward model treats these new-make differences as starting conditions: a heavier base spirit and a lighter base spirit do not follow identical trajectories even in the same cask.

For a distiller thinking in decades, the practical value is simple. A set of new-make measurements — phenolic index, basic ester and higher alcohol profile, copper contact — combined with declared cask regime and warehouse climate, yields a projected corridor for colour and chemical maturity at 10, 15, or 20 years. The model does not tell the distiller what to like. It shows where today’s decisions are likely to land relative to known whiskies that took similar paths.

4.3 What Forward Projection Produces

The output is not a single prediction but a characterisation along multiple dimensions: projected colour development over time, chemical maturity indicators, and benchmark positioning against a reference database of existing whiskies.

Colour projection follows the same Peleg kinetic curves used in the inverse direction, but starting from known parameters rather than inferring them. If the cask type, fill number, oak species, and warehouse climate are specified, the colour trajectory can be calculated with considerably tighter confidence ranges than the inverse direction achieves — because the parameters that create ambiguity in the inverse case are given as inputs in the forward case.

Benchmark matching compares the projected chemical profile against a curated database of 2,137 reference whiskies with publicly documented cask types and age statements; coverage is weighted toward Scotch single malt and bourbon, and that bias should be borne in mind. The system identifies existing whiskies that occupy similar positions in the chemical parameter space — not as quality judgments but as reference points. “Your production regime projects to a profile similar to these documented examples” is a more useful output for a distiller than an abstract score.

4.4 Applications

For **distillers**, forward projection addresses a genuine operational challenge. A new cask fill commits capital and warehouse capacity for a decade or more. Projecting the maturation trajectory from documented physics — rather than waiting ten years to find out — supports more informed decision-making. The system doesn’t eliminate the craft judgment involved in cask selection. It provides a physics-informed baseline against which craft judgment operates.

For **independent bottlers**, the calculus is different but the need is parallel. Selecting casks for purchase from distillery inventories involves evaluating young spirit against its maturation potential. Understanding how the declared cask type, fill number, and remaining maturation time project into likely chemical outcome sharpens that evaluation.

³See, for example, Rodríguez Madrera et al. (2003), *J. Agric. Food Chem.* 51, 5709–5714. DOI: 10.1021/jf0347618.

For **investors and cask owners**, forward projection serves a due diligence function. Young stock investments are bets on future maturation outcomes. A physics-based projection of where the chemistry is heading — positioned against reference benchmarks — provides a more substantive basis for valuation than provenance documentation alone.

Important distinction: The forward direction as presented in this paper projects chemistry and colour. It does not project taste. The translation from chemistry to individual sensory experience involves additional variables — receptor sensitivity, perception thresholds, individual variation — that are outside the scope of this framework. Claims are limited to what the physics supports: chemical composition and its visual expression.

5. Bidirectional Validation

Inverse and forward aren't two unrelated tools; both apply the same underlying physics model. That structure has a useful consequence: each direction validates the other.

If the forward model projects that a specific cask regime at a given age should produce colour in a certain range, and the inverse model, given a photograph of such a whisky, infers parameters consistent with those inputs — the model is internally coherent. If either direction fails, both are called into question. The physics either works in both directions or it doesn't.

Off-the-shelf unidirectional ML models don't usually offer this kind of cross-check. A neural network trained to predict age from colour photographs can achieve high accuracy on its training distribution without encoding the causal physics. It has no forward direction to cross-validate against. The point is that a physics-based structure makes such tests straightforward — not that ML approaches are inherently inferior.

That claim has limits. Internal coherence is necessary but not sufficient. The model can be internally consistent and still wrong — if the underlying physics parameters are miscalibrated, both directions will agree on incorrect outputs. External validation against known-parameter samples is required, and that work is ahead, not behind. We haven't yet used this bidirectional symmetry to tune parameters in a systematic way; the point here is that the structure allows that kind of test, and each verified sample improves both directions simultaneously because it refines a shared model rather than two independent ones.

What systematic calibration data will determine is how far confidence ranges can be pushed. The current proof-of-concept systems demonstrate that the physics is exploitable. The precision achievable with calibrated parameters is an empirical question that this paper frames rather than answers.

6. Limitations and Honest Uncertainties

A framework that overstates its capabilities invites the scepticism it deserves. This section catalogues what the system cannot do, where confidence is genuinely limited, and which problems require methods outside this approach.

6.1 What Colour Cannot Determine

Specific distillery identity. Multiple distilleries use identical cask regimes. Colour encodes cask-spirit interaction, not the distillery name. Nothing in the visible spectrum distinguishes a Speyside malt from a Highland malt aged in the same cask type for the same duration.

Precise calendar age without cask information. The same amber can emerge from six years in first-fill Pedro Ximénez sherry or eighteen years in refill bourbon. Without metadata narrowing the cask regime, age estimates carry uncertainties of ± 5 –8 years — wide enough to be directional rather than definitive.

E150a presence from static colour alone. Spirit caramel and natural oak-derived colour overlap substantially in visible spectrum. Static measurement cannot distinguish them. The dilution behaviour principle described in Section 3 offers an additional signal, but it screens and flags — it does not definitively confirm.

Document fraud. A bottle with forged provenance documentation and colour-matched contents will present colour consistent with its false claims. Colour-physics analysis verifies whether stated parameters are consistent with observed colour. It cannot detect forgery that gets the chemistry right.

Blends and complex vattings. Most commercial whisky is blended from multiple cask types, fill numbers, and ages. The resulting colour is a composite. The current framework isn't designed to deconvolve that mixture; at best it can assess whether the aggregate colour is broadly consistent with the declared maturation story. Single-cask or single-regime whiskies are the framework's strongest territory.

Caramel and filtration variability across producers. E150a use and chill filtration practices vary by producer, market, and expression — and those practices aren't always transparent. In markets or brands where caramel addition is routine and undeclared, colour-based inference is inherently less reliable. The framework acknowledges this honestly rather than pretending the variable doesn't exist.

6.2 Where Confidence Ranges Are Wide

Unusual cask types and complex finishes. Multi-stage finishing regimes (bourbon cask, transferred to port pipe, finished in rum cask) produce colour signatures outside standard single-cask models. The framework's confidence drops appropriately for these cases, but the drop is real — not a conservative disclaimer masking hidden capability.

Lighting and capture conditions. Smartphone cameras vary in colour accuracy. Ambient

lighting shifts colour temperature. Glass thickness and colour affect the liquid’s appearance. In practice, auction rooms use a mix of glassware and bottle photography; tinted glass, fill height, and meniscus shape all nudge the colour measurement. The current implementation assumes a standard clear glass and controlled framing, which is a limitation. The system normalises for lighting variables, but normalisation introduces its own error margins.

Edge cases in the reference database. The benchmark matching in forward projection relies on a reference database. Unusual production methods not well-represented in that database will produce less reliable matches. The system knows what it knows, and its confidence should reflect the boundaries of its reference data.

6.3 Where Laboratory Methods Remain Necessary

For definitive age verification, radiocarbon dating (Cook et al., 2020; DOI: 10.1017/RDC.2019.153) remains the gold standard. For precise chemical composition, gas chromatography–mass spectrometry provides resolution that colour analysis cannot approach. For legal or regulatory purposes, accredited laboratory analysis carries evidentiary weight that computational inference does not.

The framework’s position is complementary. It provides a rapid, non-destructive, accessible screening layer. It does not replace the laboratory. It occupies a space in the assessment pipeline where the laboratory is too expensive, too slow, or too destructive for routine use. Any deployment in regulated contexts — retail or auction — would need to be framed explicitly as screening, not certification.

6.4 The Calibration Horizon

Current confidence ranges are based on published literature parameters and proof-of-concept testing. Systematic calibration — known-age, known-cask samples photographed under controlled conditions — will tighten or fail to tighten these ranges. The framework is designed to be calibrated, not to operate indefinitely on literature-derived parameters alone.

Our working assumption is that there’s practical value if age ranges can be narrowed to within a few years, given known cask type. Whether that’s achievable — and how different segments of the industry value it — is an empirical question. If calibration data show that colour-physics inference doesn’t tighten meaningfully beyond ± 7 years regardless of metadata, the framework needs either better input data or honest repositioning. We do not know which outcome awaits. The paper exists to frame the question and demonstrate that the physics is exploitable. The calibration data will answer how far that exploitation reaches.

6.5 Provenance Modifiers, Not Oracles

Colour-physics alone cannot tell you where a whisky was made. What it can do, in combination with climate and historical context, is rule out stories that are physically implausible.

Regional evaporation and extraction data are clear: a five-year cask in Taiwan or India can show an extraction profile comparable to 12–15 years in a cool Scottish dunnage warehouse,

and three to four times the annual angel's share.⁴ Historical practice adds another layer: large-scale cask finishing is a post-1980 phenomenon, and widespread use of E150a to standardise colour is largely a post-1970 behaviour codified in the 2009 Scotch Whisky Regulations.

The framework uses these as modifiers. It will not declare that a given cask "must be" Scottish or Indian. It will flag when a claimed vintage predates the emergence of sherry finishing as a commercial practice, or when a supposed 1960s bourbon-matured malt shows the colour signature of modern caramel standardisation.

⁴See climate-maturation data in Jiménez-Moreno et al. (2019), *Food Res. Int.* 125, 108577. DOI: 10.1016/j.foodres.2019.108577, and Trillo-Ollero et al. (2024), *Food Chem.* 140128. DOI: 10.1016/j.foodchem.2024.140128.

7. Possibilities

This paper has presented a physics framework and two proof-of-concept systems. What follows is not a market projection. It is an observation about what becomes available to the industry if the framework calibrates successfully.

For collectors considering high-value purchases, a non-destructive screening layer changes the risk calculus. A £50,000 bottle that passes colour-physics consistency checks for its stated cask type and age isn't authenticated — but a bottle that *fails* those checks warrants further investigation before the cheque is written. Screening is not certification. It is the difference between purchasing blind and purchasing with one additional data point.

For precision distillers and craft producers, forward projection offers something that hasn't existed: a physics-informed preview of where production decisions lead. The whisky industry operates on the longest feedback loops in food production. A distiller making cask-fill decisions today won't taste the consequences for a decade or more. Reducing that informational gap — even partially — has practical value for operations measured in decades and investments measured in millions.

For auction houses processing hundreds of lots, a rapid consistency screen could flag the 10% of lots worth sending for laboratory analysis. Current practice is either test everything (prohibitively expensive) or test nothing at the liquid level (the status quo). A middle path — screen everything, test the flagged outliers — becomes possible with a non-destructive, photograph-based approach.

For the whisky industry broadly, the framework introduces a concept that doesn't currently have a name: liquid-level intelligence. The industry has label intelligence (provenance documentation), packaging intelligence (authentication features), and palate intelligence (expert assessment). What it lacks is a systematic, physics-based approach to the liquid itself. Whether this framework or some evolution of it fills that gap, the gap exists and the physics to address it is published.

None of these possibilities depend on the framework being perfect. They depend on it being useful — providing information that currently isn't available, at a cost and speed that make routine use practical. The calibration data will determine how useful.

What comes next is evidence.

References

Oak Chemistry and Colour

- Vivas N, Bourden-Nonier MF, Vivas de Gaulejac N, et al. (2020). Phenolic extraction and color in new oak cask aging. *J. Inst. Brew.*, 126, 262–270. DOI: 10.1002/jib.586
- Nocera A, Guerrero-Chanivet M, García-Moreno MV, et al. (2023). Oak species, toasting, and aging effects on phenolic composition. *J. Agric. Food Chem.*, 71(4), 2158–2168. DOI: 10.1021/acs.jafc.3c00501. PMID: PMC10835728
- Żyźelewicz D, et al. (2023). Wood compounds and quality review. *Foods*, 12(1), 144. DOI: 10.3390/foods12010144. PMID: PMC9866382
- Collins TS, et al. (2016). Spectral-phenolic correlation in 68 whiskies. *J. Sci. Food Agric.*, 96(13), 4557–4567. DOI: 10.1002/jsfa.6960
- Delgado-González MJ, et al. (2021). Colour and phenolic evolution kinetics in oak-aged spirits. *Food Chem.*, 360, 130001. DOI: 10.1016/j.foodchem.2021.130001
- MacKenzie WM, Aylott RI (2004). Analytical strategies to confirm Scotch whisky authenticity. Part II: mobile brand authentication. *Analyst*, 129, 607–612. DOI: 10.1039/b403068k
- (2021). Color evolution kinetics in CIE L*a*b*. *Food Chemistry*, 340, 128122. DOI: 10.1016/j.foodchem.2020.128122

Extraction Kinetics

- Viriot C, Scalbert A, Lapierre C, Moutounet M (1993). Ellagitannins and lignins in aging of spirits in oak barrels. *J. Agric. Food Chem.*, 41(11), 1872–1879. DOI: 10.1021/jf00035a013
- Delgado-González MJ, et al. (2024). Fill number and Peleg kinetics in seasoned oak. *Foods*, 13(20), 3279. DOI: 10.3390/foods13203279. PMID: PMC11827016
- Rodríguez Madrera R, Suárez Valles B, Picinelli Lobo A (2003). Chemical and sensory changes in fresh cider spirits during oak aging. *J. Agric. Food Chem.*, 51(19), 5709–5714. DOI: 10.1021/jf0347618
- (2021). Peleg model for color-time kinetics. *Food Control*, 125, 107934. DOI: 10.1016/j.foodcont.2021.107934
- Jordão AM, et al. (2008). Ethanol concentration effect on oak extraction. *Food Chemistry*, 106(1), 381–387. DOI: 10.1016/j.foodchem.2007.12.054
- Psarra E, Makris DP, Kallithraka S (2011). Temperature effect on extraction kinetics. *Chem. Eng. Res. Des.*, 89(6), 617–623. DOI: 10.1016/j.cherd.2010.11.003
- Trillo-Ollero A, et al. (2024). Cask size and surface-to-volume ratio effects on extraction kinetics in spirits. *Food Chem.*, 140128. DOI: 10.1016/j.foodchem.2024.140128
- Jiménez-Moreno N, et al. (2019). Temperature acceleration of phenolic extraction from oak into wine and spirits. *Food Res. Int.*, 125, 108577. DOI: 10.1016/j.foodres.2019.108577

Cask Type Differentiation

- Valcárcel-Muñoz MJ, et al. (2021). Sherry vs new oak direct comparison in brandy. *Foods*, 10(2), 264. DOI: 10.3390/foods10020264. PMID: PMC7911483

Roullier-Gall C, Sigrist B, Bruchet A, et al. (2018). FT-ICR-MS bourbon vs sherry whisky molecular fingerprints. *Front. Chem.*, 6, 29. DOI: 10.3389/fchem.2018.00029

Nocera A, et al. (2020). Micro-oxygenation and phenolic polymerisation. *OENO One*, 54(3). DOI: 10.20870/oeno-one.2020.54.3.3114

Fill Number Quantification

Puech JL, Rabier P, Bories-Azeau J, Sarni F, Moutounet M (1990). Determination of ellagitannins in extracts of oak wood and in distilled beverages matured in oak barrels. *J. Assoc. Off. Anal. Chem.*, 73(4), 498–501. PMID: 2211469

Chira K, Teissedre PL (2015). Extraction of oak volatile and non-volatile compounds and sensory profile of wines aged with French oak chips. *Food Chemistry*, 69, 90–97. DOI: 10.1016/j.foodchem.2013.01.039

Canas S, et al. (2022). Successive extraction of ellagitannins from chestnut oak staves in wine spirits. *LWT — Food Sci. Technol.*, 158, 113664. DOI: 10.1016/j.lwt.2022.113664

Valcárcel-Muñoz MJ, et al. (2021). Phenolic development in used vs new casks. *Foods*, 10(7), 1594. DOI: 10.3390/foods10071594

Sensory–Chemistry Relationships

Lee KY, Paterson A, Piggott JR (2000). Perception of whisky flavour reference compounds by Scottish distillers. *J. Inst. Brew.*, 106(4), 203–208. DOI: 10.1002/j.2050-0416.2000.tb00068.x

Soares S, Brandão E, Mateus N, de Freitas V (2020). Tannin–salivary protein binding and astringency mechanism. *Molecules*, 25(11), 2590. DOI: 10.3390/molecules25112590. PMID: PMC7321337

Haug L, et al. (2023). Whisky sensory-composition correlation. *Foods*, 12(17), 3190. DOI: 10.3390/foods12173190. PMID: PMC10556146

Czarniecka-Skubina E, et al. (2022). Flavor compounds in spirits mapping. *Molecules*, 27(21), 7273. DOI: 10.3390/molecules27217273. PMID: PMC9656916

Fujieda T, Tanaka T, Suwa Y, et al. (2008). Isolation and structure of whiskey polyphenols produced by oxidation of oak wood ellagitannins. *J. Agric. Food Chem.*, 56(16), 7305–7310. DOI: 10.1021/jf8012713. PMID: 18672883

E150a and Dilution Behaviour

Licht BH, Shaw K, Smith C, Mendoza M, et al. (1992). Characterization of caramel colours and components. *Food Chem. Toxicol.*, 30(5), 365–373. PMID: 1644378

Kew W, Goodall I, Clarke D, Sherwood J (2017). Chemical diversity and complexity of Scotch whisky as revealed by high-resolution mass spectrometry. *J. Am. Soc. Mass Spectrom.*, 28(1), 200–213. PMID: PMC5174148

Mosedale JR, Puech JL (1998). Wood maturation of distilled beverages. *Trends Food Sci. Technol.*, 9(3), 95–101. DOI: 10.1016/S0924-2244(98)00024-7

EFSA Panel on Food Additives and Nutrient Sources (2011). Scientific opinion on the re-evaluation of caramel colours (E150a-d). *EFSA Journal*, 9(3), 2004.

Fraud Documentation

- Cook GT, Dunbar E, Naysmith P, et al. (2020). A ^{14}C chronology for the world's rarest whiskies. *Radio-carbon*, 62(1), 51–62. DOI: 10.1017/RDC.2019.153
- Smith HM, et al. (2019). Stable isotope approaches for authentication of Scotch whisky. *Scientific Reports*. PMID: 31142757
- Rare Whisky 101 (2018–2020). Market analysis reports: estimated £41M suspect bottles in circulation. [Industry report, not peer-reviewed.]

Expert Reliability

- Morrot G, Brochet F, Dubourdieu D (2001). The color of odors. *Brain Lang.*, 79(2), 309–320. DOI: 10.1006/brln.2001.2495
- Quigley-McBride A, Benton CP, Johnston RA, Sherlock C (2018). Primacy and recency effects in serial tastings of whisky. *Learn. Behav.*, 46, 504–510. DOI: 10.3758/s13420-018-0346-9
- Gawel R, Godden PW (2008). Evaluation of the consistency of wine quality assessments from expert wine reviewers. *Aust. J. Grape Wine Res.*, 14(1), 1–8. DOI: 10.1111/j.1755-0238.2008.00001.x
- Ashton RH (2012). Reliability and consensus of experienced wine judges. *J. Wine Econ.*, 7(1), 77–87. DOI: 10.1017/S1931435512000099

Authentication Reviews

- Okolo CA, Kilcawley KN, O'Connor C (2023). Recent advances in whiskey analysis for authentication, discrimination, and quality control. *Compr. Rev. Food Sci. Food Saf.*, 22(6), 4957–4992. DOI: 10.1111/1541-4337.13249. PMID: 37823807
- Chen X, Sun B, Deng Q, Dong W, Sun X (2025). Current trends and perspectives on multi-analytical methodologies in spirit authentication. *Food Chemistry*, 493(Pt 3), 145898. DOI: 10.1016/j.foodchem.2025.145898. PMID: 40829444
- Zhang Q, et al. (2025). Flavor characteristics and formation mechanisms in spirits: A case study in whisky. *Food Res. Int.*, 203, 115901. DOI: 10.1016/j.foodres.2025.115901. PMID: 40022408
- de Souza CA, et al. (2025). The science of aging: Understanding phenolic and flavor development in whisky. *Foods*, 14(15), 2739. DOI: 10.3390/foods14152739. PMID: 40807675

Colour as Maturation Indicator

- Gracie K, et al. (2022). Gold nanoparticle assay for whisky maturation marker detection. *ACS Appl. Nano Mater.*, 5(10), 14379–14388. DOI: 10.1021/acsanm.2c03406. PMID: PMC9624259
- (2023). PLS regression of A420/A520 for vintage discrimination. *Sensors*, 23(2), 541. DOI: 10.3390/s23020541. PMID: PMC9857561