

Fixed-Income Securities in Discrete Time

Lorenzo Naranjo

March 2026

Bond pricing is one of the most natural applications of the stochastic discount factor framework. The SDF approach is particularly powerful here because it handles an entire term structure — bonds of all maturities — within a single, unified pricing equation. By specifying the SDF and the dynamics of a small number of state variables, we obtain closed-form or recursively computable prices for bonds of every maturity simultaneously, without having to model each maturity separately. The framework also cleanly separates the physical dynamics of interest rates from the risk-neutral dynamics that determine prices, making it transparent how risk premia shape the yield curve.

This notebook develops the theory in discrete time, starting from the basic pricing equation for zero-coupon bonds and working through three progressively richer term-structure models: the Vasicek (1977) model with constant risk premia, the Cox et al. (1985) model with conditionally heteroscedastic interest rates, and the essentially affine class of Duffee (2002). We conclude with the general N -factor affine model of Duffee and Kan (1996) and Dai and Singleton (2000). The same ideas extend naturally to continuous time, where Itô's lemma replaces the log-normal approximation — that treatment is taken up in the [bond pricing in continuous time](#) notebook.

Zero-Coupon Bonds

A *zero-coupon bond* with maturity n is a claim to one unit of the numeraire — say, one dollar — paid with certainty at time $t+n$, with no intermediate cash flows. It is the simplest possible fixed-income security, and every coupon bond can be decomposed into a portfolio of zero-coupon bonds, so understanding their pricing is the foundation of the entire term structure.

Denote the price of the n -period zero-coupon bond at time t by P_{nt} . One period later the same bond has one fewer period to maturity and trades at $P_{n-1,t+1}$, which is random from

the perspective of t . Applying the fundamental SDF pricing equation to this one-period payoff gives

$$P_{nt} = E_t(M_{t+1}P_{n-1,t+1}). \quad (1)$$

This is a recursion: the n -period price is determined by the $(n - 1)$ -period price one period ahead, weighted by the SDF. The boundary condition is $P_{0t} = 1$ for all t : a bond that has already matured pays its face value of one. Iterating (1) forward from this boundary condition, and applying the law of iterated expectations, gives

$$P_{nt} = E_t(M_{t+1}M_{t+2} \cdots M_{t+n}).$$

This expression is intuitive: the price of a multi-period bond is the expected value of all the discounting that will occur between now and maturity. A high SDF in any future period — indicating that wealth is scarce and marginal utility is high — will reduce the bond price, reflecting the opportunity cost of tying up resources.

It is convenient to work in logs throughout. Write $p_{nt} = \log P_{nt}$ and $m_{t+1} = \log M_{t+1}$. If the vector $(m_{t+1}, p_{n-1,t+1})$ is conditionally normal — as will be the case in all the models below — then the pricing equation (1) becomes exact in log form:

$$p_{nt} = E_t(m_{t+1} + p_{n-1,t+1}) + \frac{1}{2}V_t(m_{t+1} + p_{n-1,t+1}).$$

The variance term arises from Jensen's inequality: $E(e^X) = e^{E(X) + \frac{1}{2}V(X)}$ for a normal X . It is sometimes called the *convexity* correction — uncertainty about future prices raises the bond price relative to what a naive expectation would suggest, because the exponential function is convex.

The continuously compounded *yield to maturity* y_{nt} is defined by $P_{nt} = e^{-ny_{nt}}$, so $y_{nt} = -p_{nt}/n$. It is the constant per-period discount rate that reproduces the observed price. The one-period yield $y_{1t} = -p_{1t}$ is simply the current risk-free rate. Long yields are weighted

averages of expected future short rates, adjusted for risk premia and convexity — unpacking this decomposition is the main goal of the term-structure models that follow.

An investor who buys the n -period bond at t and sells it one period later earns the log return $r_{n,t+1} = p_{n-1,t+1} - p_{nt}$. This return is uncertain because the resale price $p_{n-1,t+1}$ depends on the short rate prevailing at $t + 1$. The Jensen-corrected expected excess return over the risk-free rate is the *log risk premium*:

$$E_t r_{n,t+1} - y_{1t} + \frac{1}{2} V_t r_{n,t+1} = -\text{Cov}_t(p_{n-1,t+1}, m_{t+1}).$$

This expression is the central result for bond risk premia. It says that the premium an investor earns from holding a long bond is determined entirely by the conditional covariance between the bond's resale price and the log SDF. If this covariance is negative — meaning the bond price tends to fall exactly when the SDF is high, i.e., in recessions — then the bond is risky in the sense that matters to investors, and it must compensate them with a positive expected excess return. Conversely, if the bond price rises in bad times, it functions as insurance and investors are willing to accept a lower — even negative — excess return. Understanding what drives this covariance, and in particular whether it is constant or time-varying, is the key empirical question in term-structure modeling.

The Vasicek Model

Vasicek (1977) is the simplest term-structure model that fits inside the SDF framework. It has a single state variable x_t representing the short-term interest rate, which evolves as a stationary first-order autoregression:

$$x_{t+1} = \mu + \phi x_t + \sigma \varepsilon_{t+1}, \quad \varepsilon_{t+1} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1).$$

The parameter $\phi \in (0, 1)$ controls persistence — how quickly the short rate reverts to its unconditional mean $\mu/(1 - \phi)$ after a shock. The parameter σ is the standard deviation of the innovation. Because ε_{t+1} is normal and the AR(1) is linear, x_{t+1} is conditionally normal given the current state, which will make all bond prices tractable.

The log SDF takes the form

$$m_{t+1} = -x_t - \frac{1}{2} \left(\frac{\lambda}{\sigma} \right)^2 - \left(\frac{\lambda}{\sigma} \right) \varepsilon_{t+1},$$

where λ is a single constant, the *market price of interest-rate risk*. The SDF loads on the same shock ε_{t+1} that drives the short rate. When $\lambda < 0$, the loading $-\lambda/\sigma$ is positive, so the SDF rises when ε_{t+1} is positive — that is, when rates rise unexpectedly. Long bonds lose value when rates rise (since $B_{n-1} < 0$), so they do poorly precisely when the SDF is high. Assets that lose value in high-SDF states carry systematic risk and must compensate investors with a positive risk premium. The calibration below uses $\lambda < 0$, consistent with the empirical evidence that long bonds earn positive excess returns on average.

The constant $-\frac{1}{2}(\lambda/\sigma)^2$ is a Jensen correction that keeps the model internally consistent: it ensures that $M_{t+1} = e^{m_{t+1}}$ satisfies the correct normalization $E_t(M_{t+1}R_{t+1}^f) = 1$ where $R_{t+1}^f = e^{x_t}$ is the gross risk-free return.

To see this, note that conditional on \mathcal{F}_t , the log SDF is normally distributed:

$$m_{t+1} \sim \mathcal{N} \left(-x_t - \frac{1}{2} \left(\frac{\lambda}{\sigma} \right)^2, \left(\frac{\lambda}{\sigma} \right)^2 \right).$$

Applying the log-linear pricing equation to a one-period bond — whose payoff is simply 1, so $p_{0,t+1} = 0$ — gives

$$p_{1t} = E_t(m_{t+1}) + \frac{1}{2} V_t(m_{t+1}) = -x_t - \frac{1}{2} \left(\frac{\lambda}{\sigma} \right)^2 + \frac{1}{2} \left(\frac{\lambda}{\sigma} \right)^2 = -x_t.$$

The Jensen correction from the variance exactly cancels the $-\frac{1}{2}(\lambda/\sigma)^2$ in the mean of m_{t+1} , leaving $p_{1t} = -x_t$ and therefore $y_{1t} = x_t$. The state variable is the short rate — a convenient normalization that gives x_t a direct empirical interpretation.

Solving the Model

For $n > 1$, guess the affine form

$$p_{nt} = A_n + B_n x_t.$$

Substituting into the log-linear pricing equation and matching coefficients in x_t yields the Riccati recursions:

$$\begin{aligned} A_n &= A_{n-1} + B_{n-1}(\mu - \lambda) + \frac{1}{2}B_{n-1}^2\sigma^2, \\ B_n &= \phi B_{n-1} - 1, \end{aligned} \tag{2}$$

with $A_0 = 0$ and $B_0 = 0$. The B_n recursion is linear and solves to

$$B_n = -\frac{1 - \phi^n}{1 - \phi}.$$

When $|\phi| < 1$, $B_n \rightarrow -1/(1 - \phi)$ as $n \rightarrow \infty$: long-duration bonds have exposures to the short rate that saturate at a finite level. The A_n coefficients are then determined by iterating the first recursion forward from $A_0 = 0$.

Implications

The affine solution has two notable implications. To compute the risk premium, note that $p_{n-1,t+1} = A_{n-1} + B_{n-1}x_{t+1}$ is linear in ε_{t+1} with loading $B_{n-1}\sigma$, while m_{t+1} loads on the same shock with coefficient $-\lambda/\sigma$. Their conditional covariance is therefore $\text{Cov}_t(p_{n-1,t+1}, m_{t+1}) = B_{n-1}\sigma \cdot (-\lambda/\sigma) = -B_{n-1}\lambda$, and the log risk premium is

$$E_t r_{n,t+1} - y_{1t} + \frac{1}{2}V_t r_{n,t+1} = -\text{Cov}_t(p_{n-1,t+1}, m_{t+1}) = B_{n-1}\lambda.$$

This is constant — it depends only on maturity n , not on the current state x_t . Because B_{n-1} is deterministic and λ is a fixed parameter, the risk premium every investor earns from rolling into an n -period bond is the same regardless of whether rates are high or low, rising or falling. All variation in the yield curve in Vasicek therefore reflects variation in expectations about future short rates, not time-varying risk premia. This is a clean and testable prediction, and it is strongly rejected in the data: Fama and Bliss (1987) show that forward rates have predictive power for bond excess returns beyond what expectations alone can explain.

The forward rate for the period $(t + n, t + n + 1)$, defined by $f_{nt} = p_{nt} - p_{n+1,t}$, is obtained by differencing the affine bond prices. Substituting the solution for A_n and B_n gives

$$f_{nt} = \frac{\mu - \lambda}{1 - \phi} - \left(\frac{1 - \phi^n}{1 - \phi} \right)^2 \frac{\sigma^2}{2} + \phi^n \left(x_t - \frac{\mu - \lambda}{1 - \phi} \right).$$

The forward curve has three components. The first term, $(\mu - \lambda)/(1 - \phi)$, is the risk-neutral long-run mean of the short rate — the level to which all forward rates converge at long horizons. The second term is the convexity correction, which is always negative and grows in magnitude with maturity, pulling long forward rates below the long-run mean. The third term captures the current position of the short rate relative to its long-run mean: if x_t is above $(\mu - \lambda)/(1 - \phi)$, the forward curve slopes downward; if x_t is below it, the curve slopes upward. The model can therefore generate upward-sloping, flat, inverted, or slightly humped curves, but only by varying the level of x_t — the shape is otherwise tightly constrained by the single factor.

The Risk-Neutral Measure

The SDF changes the drift of the short rate from μ to $\mu - \lambda$ without altering the volatility, giving the risk-neutral dynamics $x_{t+1} = (\mu - \lambda) + \phi x_t + \sigma \varepsilon_{t+1}^*$. Under this measure bond prices satisfy $P_{nt} = e^{-y_{1t}} E_t^*(P_{n-1,t+1})$, which yields exactly the same recursions (2). The N -factor section below works out the general change-of-measure argument in full.

Figure 1 illustrates the model's yield curve predictions under three scenarios for the current short rate x_t relative to its long-run mean of 5%. When $x_t = 2\%$, rates are below their long-run mean and the market expects them to rise, producing an upward-sloping curve whose slope

diminishes at long maturities as expectations converge to the steady state. When $x_t = 5\%$, the curve is nearly flat, with a slight upward tilt generated by the positive term premium $B_{n-1}\lambda > 0$. When $x_t = 10\%$, rates are above their long-run mean and expected to fall, inverting the curve. The Vasicek model thus captures the three most commonly observed yield curve shapes. What it cannot capture is any change in those shapes that is not driven purely by the level of x_t — rich empirical patterns such as time-varying curvature or independently moving slope are inherently beyond a single-factor model.

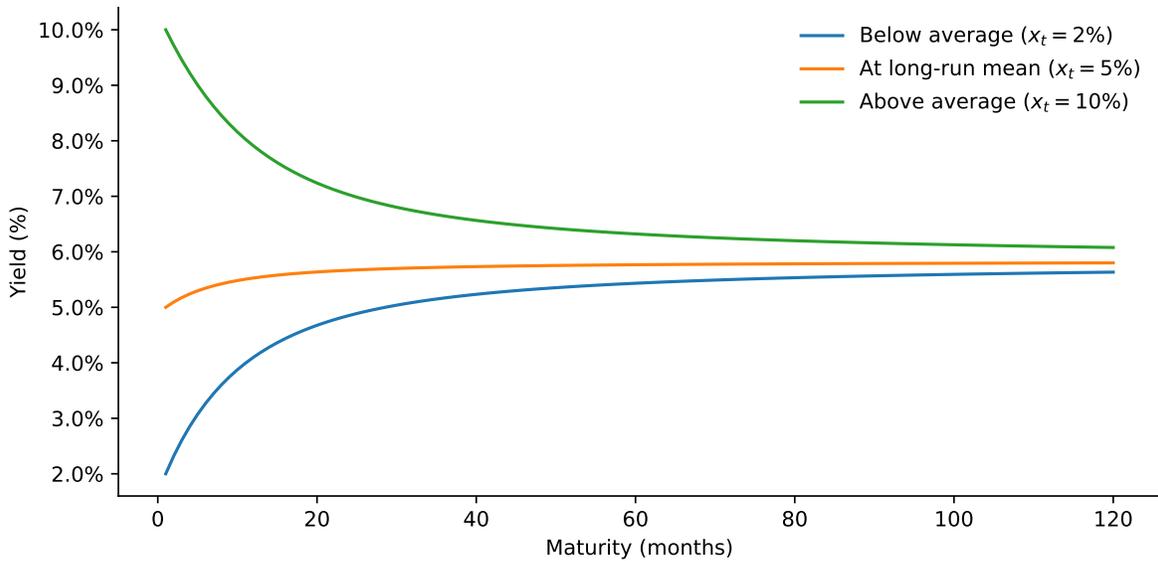


Figure 1: Yield curves in the Vasicek model for different values of the short rate x_t relative to the long-run mean.

The Cox-Ingersoll-Ross Model

Cox et al. (1985) modify the Vasicek model by making the conditional variance of the short rate proportional to its level:

$$x_{t+1} = \mu + \phi x_t + \sigma x_t^{1/2} \varepsilon_{t+1}.$$

The log SDF is adjusted accordingly:

$$m_{t+1} = -x_t - \frac{1}{2} \left(\frac{\lambda}{\sigma} \right)^2 x_t - \left(\frac{\lambda}{\sigma} \right) x_t^{1/2} \varepsilon_{t+1}.$$

Both the diffusion of the short rate and the market price of risk now scale with $x_t^{1/2}$, preserving the affine structure of log bond prices. Guessing $p_{nt} = A_n + B_n x_t$ and matching coefficients yields:

$$\begin{aligned} A_n &= A_{n-1} + B_{n-1} \mu, \\ B_n &= -1 + (\phi - \lambda) B_{n-1} + \frac{1}{2} \sigma^2 B_{n-1}^2, \end{aligned} \tag{3}$$

with $A_0 = B_0 = 0$. The B_n recursion is now *quadratic* (a Riccati equation) and generally must be iterated numerically.

The log risk premium in the CIR model is

$$E_t r_{n,t+1} - y_{1t} + \frac{1}{2} V_t r_{n,t+1} = B_{n-1} \lambda x_t.$$

Unlike Vasicek, the risk premium now inherits the time variation of x_t : when rates are high, bonds carry larger risk premia. This heteroscedastic structure is more consistent with the empirical finding that the slope of the yield curve forecasts bond excess returns. The square-root diffusion also keeps rates positive when $2\mu > \sigma^2$ — a desirable property that the Gaussian Vasicek model lacks.

A One-Factor Essentially Affine Model

Duffee (2002) proposes a class of models in which the short rate is homoscedastic (as in Vasicek) but the market price of risk is *not* affine in x_t . The essentially affine specification sets:

$$m_{t+1} = -x_t - \frac{1}{2} \left(\frac{\lambda}{\sigma} \right)^2 x_t^2 - \left(\frac{\lambda}{\sigma} \right) x_t \varepsilon_{t+1}.$$

The market price of risk $(\lambda/\sigma)x_t$ now scales with the level of the short rate, even though the short rate innovation $\sigma\varepsilon_{t+1}$ does not. Matching coefficients gives:

$$\begin{aligned} A_n &= A_{n-1} + B_{n-1}\mu + \frac{1}{2}\sigma^2 B_{n-1}^2, \\ B_n &= -1 + (\phi - \lambda)B_{n-1}, \end{aligned} \tag{4}$$

with $A_0 = B_0 = 0$. Note that B_n now satisfies a *linear* recursion (as in Vasicek, not quadratic as in CIR), so it can be solved explicitly:

$$B_n = -\frac{1 - (\phi - \lambda)^n}{1 - (\phi - \lambda)}.$$

The risk premium takes the same form as in CIR,

$$E_t r_{n,t+1} - y_{1t} + \frac{1}{2}V_t r_{n,t+1} = B_{n-1}\lambda x_t,$$

but the mechanism is different. Here conditional return volatility $B_{n-1}\sigma$ is constant — it does not depend on x_t — yet the risk premium still varies with the level of the short rate because the market price of risk $(\lambda/\sigma)x_t$ does. This separation of volatility from risk premia is the key advantage documented by Duffee (2002): the data strongly prefer models where time-varying risk premia are not tied to heteroscedastic rates.

An N -Factor Essentially Affine Model

The general affine term-structure model of Duffie and Kan (1996) and Dai and Singleton (2000) replaces the scalar short rate with a vector of state variables. Let $x_t \in \mathbb{R}^N$ follow the VAR:

$$x_{t+1} = \mu + \Phi x_t + \Sigma \varepsilon_{t+1},$$

where $\mu \in \mathbb{R}^N$, Φ is an $N \times N$ matrix, Σ is invertible, and $\varepsilon_{t+1} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, I_N)$. The log SDF is

$$m_{t+1} = -(\delta_0 + \delta_1' x_t) - \frac{1}{2} \Lambda_t' \Lambda_t - \Lambda_t' \varepsilon_{t+1},$$

$$\Lambda_t = \Sigma^{-1}(\lambda_0 + \lambda_1 x_t),$$

where $\delta_0 \in \mathbb{R}$, $\delta_1, \lambda_0 \in \mathbb{R}^N$, and λ_1 is an $N \times N$ matrix. The short rate is $y_{1t} = \delta_0 + \delta_1' x_t$. The essentially affine market price of risk Λ_t is affine in x_t , nesting the one-factor model as a special case.

Pricing the Bonds

Guessing the affine solution $p_{nt} = A_n + B_n' x_t$ with $B_n \in \mathbb{R}^N$ and applying the log-linear pricing equation gives:

$$A_n = A_{n-1} + B_{n-1}'(\mu - \lambda_0) - \delta_0 + \frac{1}{2} B_{n-1}' \Sigma \Sigma' B_{n-1},$$

$$B_n' = B_{n-1}'(\Phi - \lambda_1) - \delta_1',$$
(5)

subject to $A_0 = 0$ and $B_0 = 0$. Comparing with the one-factor case, the risk adjustment replaces $(\mu - \lambda)$ with $(\mu - \lambda_0)$ and $(\phi - \lambda)$ with $(\Phi - \lambda_1)$: the market price of risk shifts the drift of the state vector under the risk-neutral measure.

Risk-Neutral Pricing

Under the risk-neutral measure, the state dynamics become:

$$x_{t+1} = (\mu - \lambda_0) + (\Phi - \lambda_1)x_t + \Sigma \varepsilon_{t+1}^*,$$

and bond prices satisfy $P_{nt} = e^{-y_{1t}} E_t^*(P_{n-1,t+1})$. The recursion (5) is the coefficient-matching consequence of this risk-neutral pricing equation. The physical parameters (μ, Φ) govern the time-series behavior of interest rates, while the risk-neutral parameters $(\mu - \lambda_0, \Phi - \lambda_1)$ determine the cross-sectional shape of the yield curve. The essentially affine specification

allows these to differ freely, which is the key advantage documented by Duffee (2002) over completely affine models where $\lambda_1 = 0$.

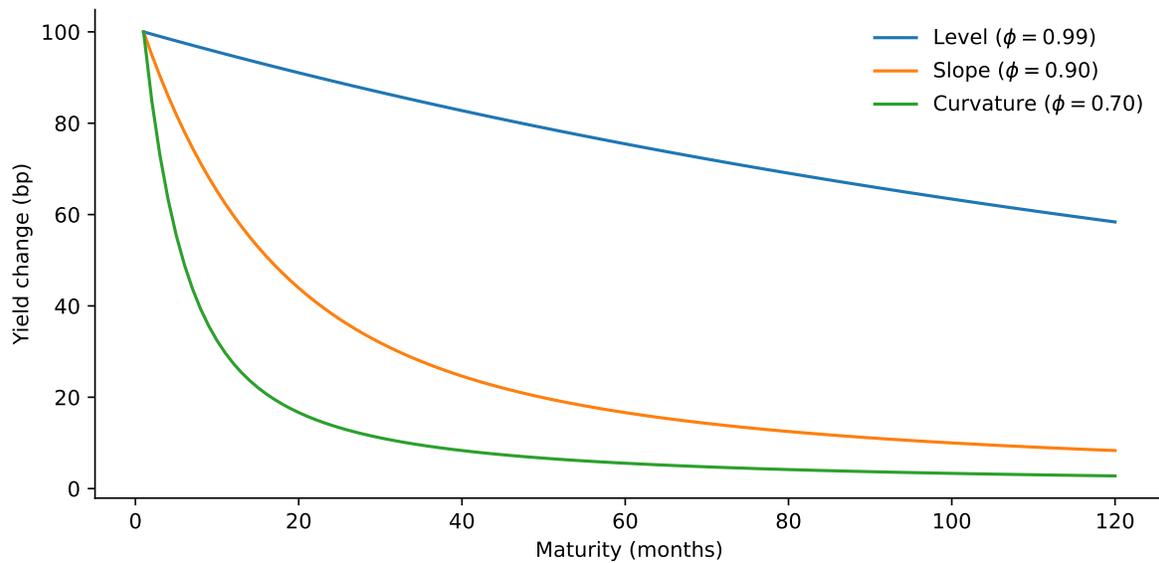


Figure 2: Yield change (in basis points) at each maturity following a 100 bp shock to each of the three factors in the essentially affine model.

Figure 2 shows how each factor transmits a 100 bp shock to yields at different maturities. All three factors affect the short rate equally — each shock raises the one-month yield by 100 bp — but they differ sharply in how quickly that effect fades with maturity. The level factor ($\phi = 0.99$) is highly persistent, so even ten-year yields rise by roughly 60 bp; the shock permeates the entire curve nearly uniformly. The slope factor ($\phi = 0.90$) mean-reverts faster, and its yield effect declines to about 8 bp by maturity 120; a positive slope-factor shock therefore tilts the curve downward, with short rates rising much more than long rates. The curvature factor ($\phi = 0.70$) mean-reverts fastest of all: the yield effect drops sharply in the first year and is negligible beyond five years, so this factor essentially only moves the shortest part of the curve. Together, these three independent patterns — level, slope, curvature — account for the bulk of yield curve variation observed in the data, a finding first documented by Litterman and Scheinkman (1991). The multi-factor affine model provides a disciplined, arbitrage-free foundation for that decomposition.

References

- Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross. 1985. "A Theory of the Term Structure of Interest Rates." *Econometrica*, 385–407.
- Dai, Qiang, and Kenneth J. Singleton. 2000. "Specification Analysis of Affine Term Structure Models." *The Journal of Finance* 55 (5): 1943–78.
- Duffee, Gregory R. 2002. "Term Premia and Interest Rate Forecasts in Affine Models." *The Journal of Finance* 57 (1): 405–43.
- Duffie, Darrell, and Rui Kan. 1996. "A Yield-Factor Model of Interest Rates." *Mathematical Finance* 6 (4): 379–406.
- Fama, Eugene F., and Robert R. Bliss. 1987. "The Information in Long-Maturity Forward Rates." *American Economic Review* 77 (4): 680–92.
- Litterman, Robert, and José Scheinkman. 1991. "Common Factors Affecting Bond Returns." *Journal of Fixed Income* 1 (1): 54–61.
- Vasicek, Oldrich. 1977. "An Equilibrium Characterization of the Term Structure." *Journal of Financial Economics* 5 (2): 177–88.