Martin Larsson and Johannes Ruf's contribution to the Discussion of 'Estimating means of bounded random variables by betting' by Ian Waudby-Smith and Aaditya Ramdas

Martin Larsson[1] Johannes Ruf[2]

[1] Department of Mathematical Sciences, Carnegie Mellon University
[2] Department of Mathematics, London School of Economics

larsson@cmu.edu j.ruf@lse.ac.uk

We congratulate Ian Waudby-Smith and Aaditya Ramdas on their comprehensive and insightful paper.

The authors construct time-uniform confidence sequences for the mean of a sequence of $[0,1]$-valued random variables $X_1, X_2, \dots$ that all have the same conditional mean, assumed to be deterministic. Specifically, fix $m \in (0,1)$ and define $\mathcal{P}^m$ as the set of probability measures on the canonical sequence space under which, for each $t \in \mathbb{N}$, the conditional expectation of X_t given the history $X_1, \dots, X_{t-1}$ is equal to the deterministic number m. The authors construct nonnegative processes (K_t) that satisfy $K_0 = 1$ and are $\mathcal{P}^m$-martingales, i.e. $\mathbb{P}$-martingales for each $\mathbb{P} \in \mathcal{P}^m$. These martingales are then used to construct anytime-valid statistical tests that in turn can be transformed into confidence sequences (see also Ramdas et al. (2022)).

In keeping with the game-theoretic probability literature, the authors refer to the processes (K_t) as *capital processes*. Let us ponder this terminology, starting with the following simple but interesting observation made by the authors: every nonnegative $\mathcal{P}^m$-martingale (K_t) with $K_0 = 1$ is of the form

$$K_t = \prod_{s=1}^{t} (1 + \lambda_s (X_s - m)) \tag{1}$$

for some predictable process (λ_t) with values in $[-(1-m)^{-1}, m^{-1}]$. *Predictable* means that each λ_t only depends on $X_1, \dots, X_{t-1}$. One may interpret λ_t as the proportion of one's capital K_{t-1} that is invested in an asset with return $X_t - m$, keeping whatever is left over 'in the pocket.' The fact that λ_t can be greater than one, or negative, poses no issue as this simply means that one may *borrow* cash to purchase more of the asset than one could otherwise afford, or *sell the asset short* to generate additional cash income. Crucially, one's capital must always remain nonnegative.

The upshot is this: not only is (K_t) the capital process produced by repeated betting; thanks to the representation (1) there is *always* an explicit trading strategy, operating on one single asset, that generates the capital process. Indeed, one has $\lambda_t = (K_t/K_{t-1} - 1)/(X_t - m)$. Given any particular (K_t) of interest, we believe insight can be gained by computing the associated trading strategy. For example, the 'diversified Kelly' capital process considered by the authors is

$$K_t^{\mathrm{dKelly}} = \frac{1}{D} \sum_{d=1}^{D} \prod_{s=1}^{t} (1 + \lambda_s^d (X_s - m)),$$

built from D separate strategies (λ_t^d), $d = 1, \ldots, D$. This is equivalent to the single strategy

$$\lambda_t = \frac{\sum_{d=1}^{D} K_{t-1}^d \lambda_t^d}{\sum_{d=1}^{D} K_{t-1}^d},$$

where (K_t^d) is the capital process generated by the d-th strategy. In other words, diversified Kelly arises from executing the capital-weighted average of the given strategies. This links it to Cover's universal portfolios (Cover (1991)).

Many of the capital processes proposed in the paper are specified in terms of a strategy (λ_t). What we find worth emphasizing is that such a (λ_t) can *always* be found, and is likely to yield insights. Finally, let us point out that the representation (1) is of course specific to the particular structure of $\mathcal{P}^m$. An interesting question is to what extent analogous representations exist for other, more complex, statistical hypotheses.

References

T Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991.

Aaditya Ramdas, Johannes Ruf, Martin Larsson, and Wouter M. Koolen. Admissible anytime-valid sequential inference must rely on nonnegative martingales. Preprint arXiv:2009.03167, 2022.