AMERICAN ASSOCIATION OF WINE ECONOMISTS

AAWE WORKING PAPER
No. 6

RETHINKING WINE INVESTMENT IN THE UK AND AUSTRALIA

James Fogarty

Editor
Victor
Ginsburgh

September 2007

Rethinking Wine Investment in the UK and Australia

James Joseph FOGARTY

University of Western Australia, School of Economics and Commerce, 35 Stirling Highway, Crawley, W.A. 6009, Australia, Email: James.Fogarty@uwa.edu.au, Phone: +61 8 6488 2918, Fax: +61 8 6488 1016

Abstract

This article presents three arguments as to why the value of wine as an investment good has typically been understated and argues that wine investment in the UK and Australia represents a value proposition. It is argued that general all vintage wine indexes understate the return the typical investor receives; that comparisons using pre-tax returns overstate the value of standard financial assets relative to wine; and that wine investment provides value in terms of allowing portfolio risk to be reduced.

1. INTRODUCTION

Broadly speaking, the return to wine literature suggests that wine investment is not a profitable activity. See Fogarty (2006, p. 544-5) for a summary table outlining details of the return to wine literature. The work presented in this article shows that in both Australia and the UK the existing literature has understated the potential value of wine investment. It is argued here that the value of wine as an investment good has been understated for the following reasons. Firstly, as the quality of any given vintage of wine is largely known in advance, wine investors will not invest in poor quality vintages. This means a wine price index based on the return to all vintages of wine understates the return achieved by the typical wine investor. As the asset return comparisons made in the literature are typically made using the return to an all vintage wine index, these comparisons understate the return achieved by the typical wine investor. Secondly, in certain countries, such as the UK and Australia, the returns to wine are generally not taxed, while the returns to standard financial assets are taxed. As the comparisons of asset returns typically made in the literature are comparisons of pre-tax returns, for Australia and the UK these comparisons understate the value of wine investment. Thirdly, the appropriate framework for assessing the value of wine as an investment class is not a comparison of the risk-return profile of wine compared to other asset classes, but rather an assessment of the role wine can play in expanding the Markowitz (1952) Investment Efficient Frontier. Even if the return to wine is lower than standard financial assets and the risk higher, if including wine in an investment portfolio reduces portfolio risk, wine should be considered a worthwhile investment class.

This paper proceeds as follows. Section two explores and comments on some of the specific extra costs incurred by those investing in wine. Section three discusses in detail the three reasons proposed for believing the existing literature has understated the potential benefits to wine investment, and presents examples that illustrate the points made. Section four presents concluding remarks.

2. SOME WINE SPECIFIC COSTS

Those considering wine investment face a number of costs that are peculiar to wine investment. The most obvious of these costs are storage costs, but transaction costs when trading wine are also much higher than the transaction costs involved when trading standard financial assets. An assessment of the return to wine that fails to account for transaction and

storage costs will overstate the value of wine investment. The nature and magnitude of these costs are discussed below.

Wine must be kept in a cool, humid, and preferably dark location where daily temperature fluctuations are minimal. For those without a wine cellar, which is most people, there are several options: have a purpose built cellar excavated; install one or more climate and humidity controlled wine cabinets; or use and pay for commercial storage at a purpose built wine storage facility. The most cost effective approach in terms of lowest annual cost per bottle held will depend on the number of bottles stored. For a vast and comprehensive wine collection excavating a cellar will, over the longer term, prove to be the most cost effective option. For a collection of only a few bottles, the capital outlay for a cellar or wine cabinet would not be justified, and so private commercial storage would be the most cost effective option.

In Australia the average price of a bottle of investment quality wine might be around $100[1], and a reasonable level of investment might be around $50,000. Assuming the investor wants to maintain some brand diversity in their portfolio, and so does not just hold the most expensive wines, a $50,000 portfolio will consist of around 500 bottles. Commercial storage rates vary significantly between Australian capital cities. Currently, storing 40 cases of wine at the Australian Wine Exchange in Perth would cost around $800 per year and around $1,500 per year at Wine Ark in Sydney. A suitable wine cabinet would cost around $6,700[2] and construction of an underground cellar many times this amount. So, for a ten year investment horizon, and a portfolio of around 500 bottles, even in Perth where storage costs are relatively cheap, the up-front cost of a wine cabinet is likely to represent the approach that yields the lowest annual per bottle storage cost.

The average per bottle price of a portfolio of premium Bordeaux wine would be substantially higher than the average per bottle price of a portfolio of premium Australian wine. For example, one might spend up to £5,000 for a case of vintage 2005 Bordeaux from a leading Chateau such as Latour, Margaux, or Cheval-Blanc. For any level of investment, a UK wine investor is therefore likely to need much less storage space than an Australian wine investor.

---

[1] Unless otherwise indicated dollars are Australian dollars and pounds are British pounds. The average exchange rates in 2006 were $AUD 1 = $US0.76 and £1 = $US 1.84.
[2] Based on the Australian manufactured Kitchener model KA568 (capacity 568 bottles).

Let us assume an investment of approximately the same order of magnitude as our Australian investor, and consider the case of a £20,000 wine portfolio, where again the investor wants some brand diversity and so does not just hold the most expensive wines. We might therefore reasonably expect an average price of around £100 per bottle. Commercial wine storage costs in the UK are usually between £3 and £5 per case per month. With an average cost of £100 per bottle, annual commercial storage costs will therefore be at least £600 and may be as much as £1,000. A suitable wine cabinet can be purchased for around £1,500[3], which is much less than the cost of excavating a cellar. So again, provided the investment horizon is sufficiently long, for a portfolio of around 200 bottles, purchasing a wine cabinet will provide the lowest annual per bottle storage cost.

Transaction costs in large liquid markets, such as the market for shares and bonds, are very low. Transaction costs for wine are high. In both the UK and Australia the buyer's premium for a bottle of wine bought at auction may be as high as 15 percent, and is normally never less than 10 percent. Details on fees charged to sellers vary depending on the amount of business the seller puts through the auction house. A potential seller is however likely to be faced with a fee of at least 10 percent of the value of their property.

For wine, storage and transaction costs are substantial. For bonds and shares there are no storage costs, and transaction costs are minimal. Failure to account for transaction and storage costs when calculating the return to wine will result in an overstatement of the return to wine. Yet despite these extra costs, this paper argues that wine can still represent a valuable addition to an investment portfolio. The reasons for this are explored in detail in the next section.

---

[3] Based on a Transtherm Ermitage M (capacity 219 bottles).

3. THREE REASONS FOR RECONSIDERING THE VALUE OF WINE INVESTMENT

Proposition One: Quality vintages can be identified in advance and people considering wine investment hold only quality vintages.

Although work presented in relatively recent times has formalised the relationship between weather and wine quality (Ashenfelter et al., 1995; Byron and Ashenfelter, 1995; De Vittorio and Ginsburgh, 1996) in a way that economics professionals can understand, the basic weather conditions that lead to good wine have long been known by viticulturalists. See for example Gladstones (1992) for historical, technical, and statistical details on wine and climatic conditions. Yet even those who have never picked up a viticulture textbook or a wine economics paper know wine quality varies between vintages. It is, for example, inconceivable that anyone in the UK interested in wine investment would not be aware 2005 was a great vintage in Bordeaux, just as it is equally impossible for someone in Australia to not know 1998 was a great vintage in Australia's Barossa Valley and Coonawarra region.

The popular press do report on exceptional vintages, but more generally, those interested in wine understand that quality varies between vintages, and because they know quality varies each year, they use vintage charts to inform their wine purchasing decisions. In particular, if the wine is from Bordeaux, as the majority of investment quality wine traded in the UK is, it is almost inconceivable that a potential wine investor would not consult one of Robert Parker's vintage charts.

While a notable critic of the econometric approach to explaining vintage variation, Parker does give numerical scores to each vintage from the key wine producing regions of France. Specifically, he gives a rating out of 100 for each vintage from each region. In the Parker rating system, if the vintage score is 96 or higher the wines are said to be extraordinary; 90 to 95 and the wines are the finest; 89 to 80 means the wines are above average; 70 to 79 means the wines are average; 60 to 69 means the wines are below average; and below 60 means the wines are appalling. Figure 1 provides a 20 year snapshot of the Parker vintage ratings for St. Emilion and Pomerol. From the details shown it is clear the quality of vintages -- as assessed by Parker -- varies significantly from year to year. Charts for the other areas of France display similar characteristics.

FIGURE 1. PARKER POINTS FOR SELECT BORDEAUX REGIONS

Source: www.erobertparker.com [accessed: 10 January 2007]
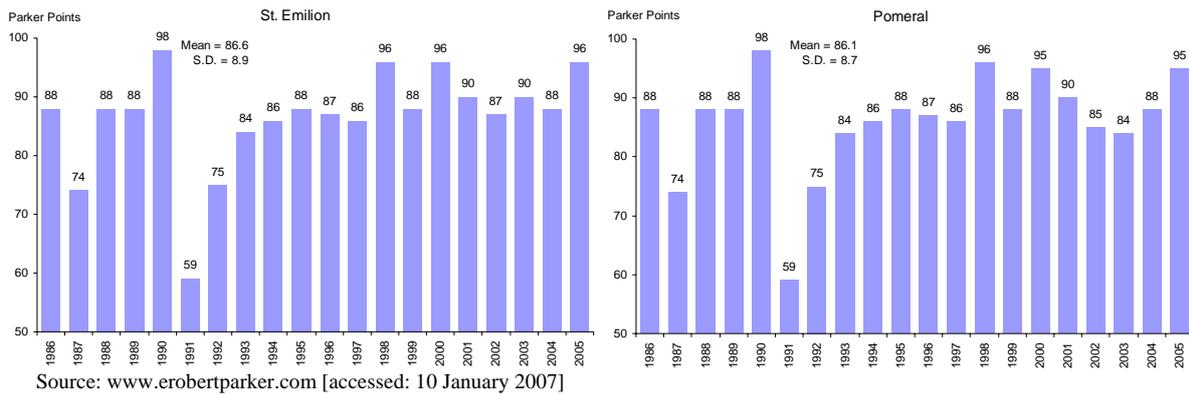
For Australian wine there are also vintage charts and wine ratings such as those of: Langton's, the largest Australian wine auction house; Jeremy Oliver, author of the Australia Wine Annual; and James Halliday, author of the Australian Wine Companion. Australian investors undoubtedly also consult a vintage chart before making the decision to invest in a particular vintage. Although no one critic has quite the same status as Robert Parker, the ratings provided by Australian critics are still highly regarded. Figure 2 shows plots of the Langton's vintage ratings for two famous Australian wine variety region combinations, Coonawarra Cabernet Sauvignon and Barossa Valley Shiraz. In the Langton's rating system the highest rating a vintage can receive is ten.

FIGURE 2. SELECT VINTAGE-VARIETY RATINGS FOR AUSTRALIA

Source: www.langtons.com.au [accessed: 10 January 2007]
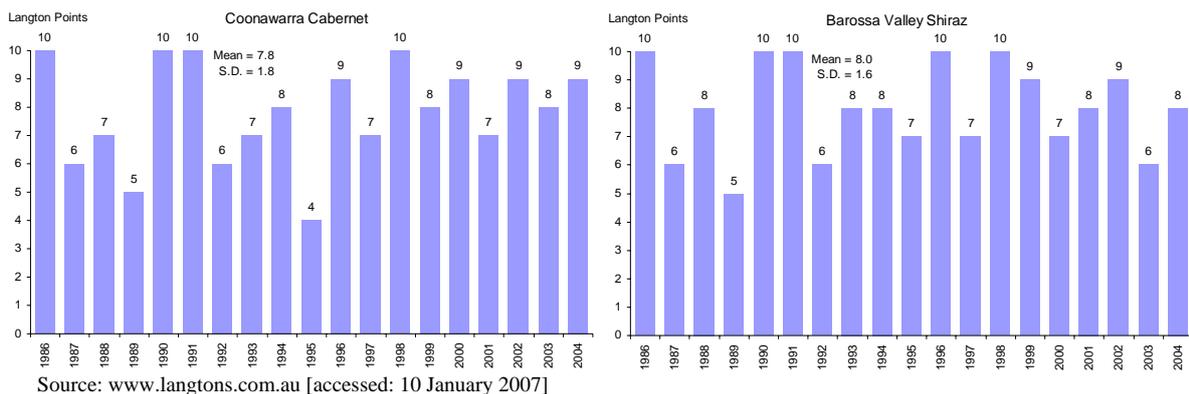
The existence of vintage charts alone does not justify moving away from a general wine price index to one based on select vintages. It is only reasonable to move to a select vintage wine index because vintage chart information is published substantially in advance of when the wine is available for sale. In this sense the equity market and the wine market are

fundamentally different. When a company lists on the stock exchange there is much uncertainty about its future profitability. Whether the management team have the skills to generate above average results is largely unknown. When a wine first appears in the secondary market there is very little uncertainty about whether the wine is above average, below average, or merely average. The quality of the vintage is largely determined by the weather conditions prevailing that year. As fine wine is sold some years after the grapes have been harvested, the wine bottled, and the vintage charts published, the quality of any given vintage is largely known prior to purchase. All but the most determinedly ignorant of wine investors will not invest in bad vintages. As such, in the case of wine, the return to a general wine index understates the return the average investor is likely to receive.

It is worth noting that there is no wine investment advice in published or on-line sources that suggests holding an investment portfolio of all wine vintages. Wine investment decisions and or recommendations, are always discussed with respect to each particular vintage. This feature of the wine market appears not to be reflected in the return to wine literature.

For shares and bonds, it is however appropriate to use, for comparison purposes, the return to an appropriately defined broad market total return index. Ensuring above average performance over the longer term, for most investors, is simply not possible. Were it possible to consistently achieve above average returns, active fund managers would consistently out-perform the market. Yet as Sharpe (1991, p. 7) explains, the return to the market is the weighted average of the return to the active and passive segments of the market, and as the passive return equals the market return, the average return to an actively managed fund must also equal the market return. Empirical work generally supports the idea that the return to active fund management, once fees are taken into consideration, is no better than the market return, and may be noticeably worse. (Frino and Gallagher, 2001; Gallagher and Jarnecic, 2002).

Proposition Two: Investors are concerned with after-tax returns not pre-tax returns. In both the UK and Australia, the returns to wine are generally not taxed. Comparisons made using pre-tax returns will therefore understate the attractiveness of wine to potential investors.

In Australia, if someone is deemed to be trading wine as an on-going dedicated business activity, the income derived from capital gains on the sale of wine are treated as income and

are taxable. There are, however, few dedicated wine investment companies in Australia, and even fewer people who trade wine as their main occupation. For most Australians the capital gains from the sale of wine will fall under the collectables exemption to capital gains tax, and so profits on the sale of wine are not taxed. This is not the case for the return to shares and bonds. In Australia dividend payments generally attract franking credits that, depending on the marginal tax rate faced by the investor, offset part or all of the individual's tax liability. Capital gains on share holdings are, however, subject to capital gains tax. There are no tax concessions associated with holding bonds, although actual tax rates will vary according to the proportion of the returns that are capital gains and the proportion that represent an income stream.

Since at least 2003, the Australian Stock Exchange has published the ASX Investment Sector Performance Report, which details the impact of taxation on asset returns. The calculations presented in each year's report show the effect of tax on investment returns for Australian residents faced with the lowest and highest marginal income tax rates, and cover a variety of asset classes. The left-hand frame of Figure 3 reproduces details for the rolling 10-year returns to December 2002, 2003, and 2004 for both government bonds and Australian shares. By considering the first set of columns in the figure it can be seen that the pre-tax, but after-expense, annual nominal return to Australian shares for the 10-year period ending December 2002 was 10.5 percent. For someone faced with the lowest marginal tax rate, the after-tax return is the same as the pre-tax return because the franking credits fully off-set the individual's tax liability. For someone faced with the top marginal tax rate, the return falls from 10.5 percent to 8.3 percent. For Australian government bonds, as shown in the second set of columns, the pre-tax, but after-expense, annual nominal return for the 10-year period ending December 2002 was 8.1 percent. For someone faced with the lowest marginal tax rate, the after-tax return falls to 6.3 percent, while for someone faced with the top marginal tax rate, the return falls to 4.1 percent. The results for the rolling 10-year periods ending December 2003 and 2004 are broadly similar. It seems reasonable to treat investment income as marginal income and in Australia, for the period under consideration, the top marginal tax rate applied to income above $50,000. It is therefore likely that potential wine investors faced the top marginal tax rate.[4]

---

[4] There have since been substantial reforms to the Australian taxation system. In particular, marginal income tax thresholds have been raised substantially.

FIGURE 3. THE EFFECT OF TAX ON ASSET RETURNS

Source: Australian Data: ASX (various years), UK Data: Armitage (2004)
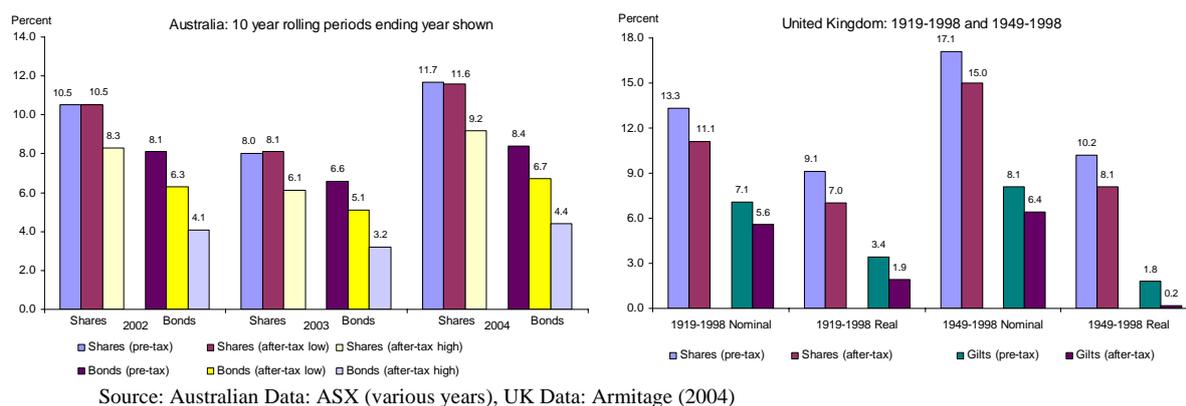
The situation in the UK is similar. If deemed to be primarily engaged in the buying and selling of wine for profit, it is possible the capital gains from trading wine could be treated as income by the taxation authorities. It is however generally the case that the capital gains to wine investment are not taxed. The returns to gilts and shares in the UK are however subject to tax. The calculations presented in Armitage (2004, p. 17) for the period 1919-1998 and 1949-1998 are helpful in understanding the impact of taxation on UK investment returns. Select key findings from the study are reproduced in the right-hand frame of Figure 3. In Armitage (2004) income from investments is treated as marginal income, and as shown in the fifth and sixth sets of columns in right-hand frame of Figure 3, for the period 1949-1998, taxation, on average, reduced the annual nominal return to UK shares from 17.1 percent to 15.0 percent, and for gilts reduced the annual average nominal return from 8.1 percent to 6.4 percent.

Although investors are primarily concerned with after-tax returns, discussions on investment returns generally take place in terms of pre-tax returns. That the tax circumstances of individuals vary markedly is the likely reason discussion focuses primarily on pre-tax returns. Yet, tax affects the return to each asset class differently. In Australia, on average, over the rolling 10-year periods ending 2002, 2003, and 2004, for those faced with the highest marginal tax rate, taxation reduced the nominal return to equities by 21.9 percent, and the nominal return to government bonds by 49.5 percent. In the UK, for the period 1949-1998, taxation reduced the average nominal return to shares by 12.3 percent and the average nominal return to gilts by 28.4 percent.

With respect to standard financial assets, taxation substantially reduces the return received by the investor. In the UK and Australia the return to wine is generally not taxed. As such, it is important comparisons of the risk-return profile of wine, compared to standard financial assets, be made using after-tax returns. Comparisons made using pre-tax returns systematically understate the value of wine investment.

For the UK, Burton and Jacobsen (2001) publish return information for red Bordeaux wine, vintage 1960 onwards, for the period between the second half of 1986 and the second half of 1996. The return estimates presented were calculated using a repeat-sales index approach. The paper provides details on the return to all Bordeaux wine, first growth Bordeaux wine, Bordeaux vintage 1961 wine, and Bordeaux vintage 1982 wine. The paper also reports details for other financial assets. However, in the case of shares, the paper does not take into consideration the effect of dividends. It is simply noted that dividends tip the balance further in favour of shares and against wine. While adjustments for the specific costs associated with trading and storing wine are made in the paper, no adjustments for the effect of taxation are made. Wine is therefore shown in a less positive light than would otherwise be the case.

The left-hand panel of Table 1 uses the wine return information published in Burton and Jacobsen and attempts to make the appropriate adjustments for a person living in the UK, so that a like-for-like comparison of asset returns can be made. Some costs, such as storage costs for wine, and transactions costs for standard financial assets, vary with the size of the portfolio. The size of the portfolio therefore matters. For the illustrative case shown in Table 1, it is assumed an investment of £20,000 is made in each asset class. In the case of wine it has been assumed the buyer's premium is 15 percent, the seller's fee is 10 percent and that as discussed previously, a wine cabinet requiring an initial investment of £1,500 is used to store the wine. For the standard financial assets the transaction costs are based on the rates advertised by Barclays, and the returns are total returns. The tax rates used are the representative effective rates in Armitage (2004), for the period 1949-98, of 12.3 percent for shares and 28.4 percent for gilts. The all vintage return details are based on the details shown in Burton and Jacobsen (2001) under the heading all wine, and the top vintages return represents the arithmetic mean return of the values reported for vintage 1961 and vintage 1982 wine.

Interestingly, and as can be seen by comparing the headline pound return and the post adjustment return information shown in Table 1, the adjustments do not change the relative ordering of the assets in terms of returns. Both before-adjustments and after-adjustments shares provide the highest returns, followed by the return to a portfolio of vintage 1961 and 1982 wines, then gilts, and lastly a portfolio of all wine vintages. Specifically, the adjustments for transaction costs, storage costs, and taxation result in the annual nominal return to a portfolio of all vintages of wine falling from 12.4 percent to 5.8 percent; the annual nominal return to gilts falling from 14.5 percent to 8.9 percent; the annual nominal return to a portfolio of vintage 1961 and 1982 wine falling from 15.6 percent to 10.5 percent; and the annual nominal return to UK shares falling from 15.6 percent to 10.5 percent.

Fogarty (2006) reports information on the return to high quality Australian wines, regardless of variety, vintage 1965 onwards, for the period from the last quarter of 1989 to the last quarter of 2000. The approach used in the paper is an adjacent period hedonic regression approach, and return information is provided for an all wine portfolio as well as for a portfolio of expensive wine, and a portfolio of less expensive wine. The paper does not present return information for select known high quality wine vintages. In the paper return comparisons are made between an all vintage wine index, a total return Australian share index and three month treasury bills. Again shares are noted as the investment class that provide the highest returns. In the paper the results reported for wine are not adjusted for storage and transaction costs, and the results reported for shares and treasury bills are not adjusted for the effect of tax. As such it is unclear whether the paper understates or overstates the return to wine investment relative to shares and treasury bills.

As a central proposition of this paper is that wine investors only invest in high quality vintages, it is important to consider not just the return to an all vintage Australia wine index but the return to a select high quality vintage wine index. It has been possible to obtain access to the data set used in Fogarty (2006) and re-estimate the return to wine using the repeat-sales approach. The relative merits of the hedonic approach versus the repeat-sales approach, or even a hybrid approach are not discussed here. Rather, the repeat-sales approach has been used because it allows estimates of the return to select vintages of high quality wine to be obtained while the adjacent period hedonic approach, due to its high degrees of freedom requirement, does not.

The right-hand panel of Table 1 attempts to make the appropriate adjustments for a person living in Australia, so that a like-for-like comparison of asset returns can be made. For Australia it is assumed the investor is willing to invest $50,000, the buyer's premium is 15 percent, and the seller's fee is 10 percent. As discussed previously, the average price of investment quality Australian wine is much lower than that of investment quality Bordeaux wine, and so a relatively large wine cabinet costing $6,700 is required. For standard financial assets, transaction costs are based on the Commsec advertised rates for share trades. Both the share returns and the bond returns are total returns. The tax rates used represent average rates calculated from the detail reported in Figure 1 for a person facing the top marginal tax rate, and are 21.9 percent for shares, and 49.5 percent for bonds. As can be seen by considering the right-hand panel of Table 1, the results for Australia are more variable than the results observed for the UK.

Over the sample period the headline return to bonds is greatest, followed by the return to a portfolio of vintage 1986, 1990, and 1994 wine, then shares, and lastly the return to a portfolio of all vintages of wine. The effect of transaction costs, storage costs, and taxation do however change the relative ranking of the assets. After-tax and relevant costs, the best performing asset class is a portfolio of 1986, 1990 and 1994 wine, followed by shares, then bonds, and lastly a portfolio of all vintages of wine. Specifically, adjusting for costs and taxes sees the annual nominal return to a portfolio of vintage 1986, 1990, 1994 wine fall from 15.6 percent to 9.3 percent; the annual nominal return to Australian shares fall from 14.8 percent to 8.8 percent; the annual nominal return to Australian bonds fall from 15.9 percent to 7.3 percent; and the annual return to a portfolio of wine from all vintages fall from 12.9 percent to 5.7 percent[5]. The specific annualised return information for both the UK and Australia is shown in Figure 4.

---

[5] Using the hedonic return estimates in Fogarty (2006) directly, and the same assumptions regarding taxation and transaction costs used in this paper, the headline return to an all wine portfolio is estimated to be $133,688 (annual return 13.9 percent), and the after transaction and storage cost return is estimated to be $56,120 (annual return 7.1 percent). So while the hedonic approach gives a higher estimate of the return to an all vintage wine portfolio, the all vintage wine portfolio is still the worst performing asset class.

TABLE 1 ADJUSTED ASSET RETURN COMPARISON

Details	Sample Period 1987 through 1996				Sample Period 1990 through 2000			
	Bordeaux Wine		UK Investments		Australian Wine		Australian Investments	
	All Vintages	Top Vintages	Shares	Gilts	All Vintages	Top Vintages	Shares	Bonds
	£	£	£	£	$	$	$	$
Purchase price	20,000	20,000	20,000	20,000	50,000	50,000	50,000	50,000
Sale price	44,169	65,460	77,836	57,724	118,231	163,322	149,008	168,078
Headline return	24,169	45,460	57,836	37,724	68,231	113,322	99,008	118,078
Less costs								
Transaction	7,417	9,546	932	267	19,323	23,832	873	949
Storage	1,500	1,500	-	-	6,700	6,700	-	-
Taxation	-	-	7,114	10,714	-	-	21,683	58,448
Actual Return	15,252	34,414	49,790	26,743	42,208	82,790	76,452	58,680

Note: When represented this way, returns are not homogenous with respect to time. In particular, the return to shares for Australia is affected by poor performance at the start of the sample period. However, given the size of transaction costs, an approach whereby the period-by-period gain or loss is realised and transferred to or from a separate bank account, and which would be homogeneous with respect to time, is intuitively unreasonable.

FIGURE 4. ADJUSTED ANNUAL RETURNS TO VARIOUS ASSETS
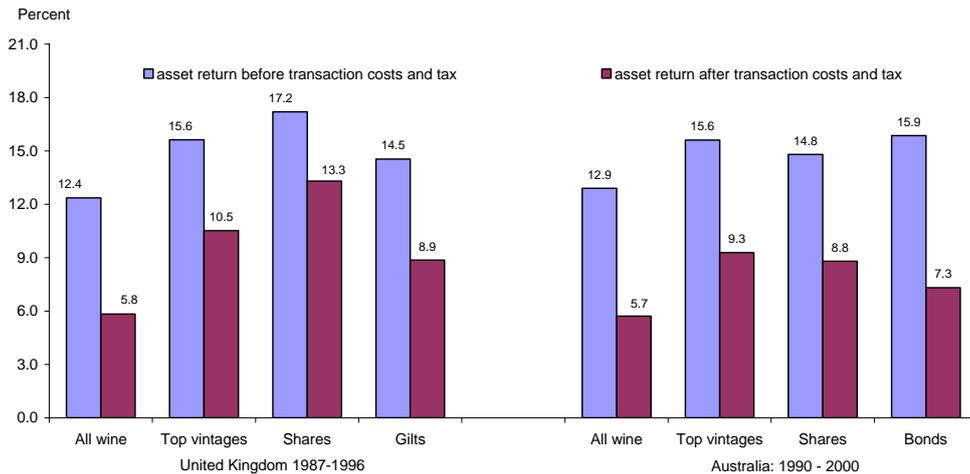
From the detail shown in Table 1 it appears, with the exception of bonds in Australia, the relative gains wine enjoys because of its tax status, are, broadly speaking, of the same order of magnitude as the extra transaction and storage costs associated with wine investment. This is an interesting result and something that would be worth exploring across other markets.

Proposition Three: The value of wine as an investment good should not be determined by a series of pair-wise risk-return comparisons across asset classes but rather by whether adding wine to an asset portfolio can expand the Markowitz (1952) Efficient Investment Frontier.

Investors are concerned about both risk and return. As such, a simple pair-wise comparison of risk and return information across asset classes is not enough to determine whether or not a particular asset class is worth holding. The expected return to a portfolio of assets depends only on the expected return to each individual asset and the proportion of funds invested in each asset. Portfolio risk on the other hand depends on the proportion of funds invested in each asset, the individual asset variances, and the correlation of asset returns.

Mean-variance analysis, as formulated by Markowitz, involves the symmetric balancing of risk and return, and is based on two fundamental propositions: (i) no matter how you define return, investors prefer higher returns to lower returns; and (ii) for any given level of return, investors prefer less risk to more risk. As such, wine may add value to a portfolio if it is an asset with high returns, low risk, or has returns that are not strongly correlated with the returns to other financial assets.

In mean-variance analysis, a portfolio is said to be efficient, if, and only if, for any given expected portfolio return, there is no alternate portfolio with lower variance. Specifically, efficient portfolios are found by solving a series of quadratic programming problems with linear constraints. Provided the covariance matrix is positive definite, there is a unique set of portfolio weights for each expected portfolio return. Once the set of quadratic programming problems are solved, the output can be plotted in two-dimensional risk-return space to describe the efficient frontier.[6] Even if the expected return to wine is lower than for shares, and even if the risk is higher, if adding wine to the portfolio mix allows risk to be lowered without compromising expected return, then wine should be considered a worthwhile investment class. The general mean-variance framework can be set out as shown below.

Let N be the number of assets under consideration, let $\boldsymbol{\mu}$ be the $N \times 1$ mean return vector, and let $\mathbf{e}$ be an $N \times 1$ vector of ones. Let $\mathbf{w}$ be an $N \times 1$ vector whose elements w_i represent the different weight allocations to the N assets, and let $\mathbf{e}'\mathbf{w} = 1$. Further, assume a market for short selling does not exist so that all the elements of $\mathbf{w}$ are non-negative. Given such

---

[6] It is worth noting there have been extensive refinements to the standard mean-variance framework since it was first proposed. Some of the more interesting improvements in the approach are based on replacing the estimated return vector or covariance matrix with a weighted average of two estimates, where the weights are determined by the data. For example see Jorion (1985) for an example of the approach applied to expected returns, and Ledoit and Wolf (2003) for a general approach to finding a shrinkage covariance matrix.

notation, the expected return to a portfolio of assets is $\mathbf{w'\mu}$. If $\mathbf{\Omega}$ is used to denote the $N \times N$ covariance matrix -- which can be estimated using mean squares and cross products -- the variance of the portfolio is $\mathbf{w'\Omega w}$. The risk of the portfolio can then be measured by its standard deviation, $\sqrt{\mathbf{w'\Omega w}}$.

The efficient frontier is found by solving the constrained optimisation problem shown at (1):

(1)
$$\text{min.} \ \mathbf{w'\Omega w}$$
$$\text{s.t.} \ \mathbf{w'\mu} \geq \mathbf{w'\mu}_L;$$
$$\mathbf{w} \geq \mathbf{0};$$
$$\mathbf{e'w} = 1.$$

In (1) $\mathbf{w'\mu}_L$ is some suitably chosen minimum portfolio return less than the maximum portfolio return achievable. Once set $\mathbf{w'\mu}_L$ is varied upwards in increments until $\mathbf{w'\mu}_L = \mathbf{w'\mu}_U$, where $\mathbf{w'\mu}_U$ is some suitably chosen upper limit to portfolio return. As long as $\mathbf{\Omega}$ is positive definite, and $\mathbf{w'\mu}_L$ and $\mathbf{w'\mu}_U$ are chosen with care, there will be a unique set of optimal portfolio weights for each level of expected portfolio return. The solution values describe the efficient investment frontier and can be plotted in two dimensional risk-return space.

Adjusting for taxation shifts any rate of return comparison in favour of the return to wine. Taking into consideration transaction and storage costs shifts any rate of return comparison in favour of standard financial assets. Individual circumstances are different, and so the effect of taxation and storage costs differ across individuals. Nevertheless, it is worth attempting to make the appropriate adjustments and in so doing obtain an approximate standardized asset return comparison. The discussion and results presented below are necessarily somewhat stylised and so should be seen as indicative only. The results do, however, serve to illustrate a valuable point regarding wine investment. Namely, that wine can add value to an investment portfolio not only from the return it provides, but also from its ability to provide portfolio risk diversification.

Given the limited number of observations for the UK, it is really only possible to illustrate the mean-variance approach using the Australian data. Although, as there are only 44 quarterly observations and as the example considers five assets, even for the Australian data, the covariance matrix is not estimated with great precision. Figure 5 below plots the efficient

investment frontier based on returns adjusted for tax, transaction costs, and storage costs, under two sets of assumptions.

The efficient frontier plots shown in the left-hand panel of Figure 5 use the following assumptions. To reflect transaction costs and taxation effects, the nominal quarterly return to Australian shares was reduced by 22 percent and the nominal quarterly return to Australian bonds was reduced by 50 percent. To reflect the effect of storage and transaction costs the quarterly nominal return to a select portfolio of vintage 1986, 1990, and 1994 wine was reduced by 30 percent. The quarterly returns to US shares are unhedged Australian dollar total returns, and to calculate the effect of taxation it has been assumed that 70 percent of the return represents capital growth, and 30 percent relates to dividend payments. As such, the nominal returns were reduced by 32 percent to account for the effect of taxation and transaction costs. For US bonds the returns are unhedged Australian dollar total returns and it has been assumed that 70 percent of the return represents income and 30 percent reflects capital growth. The nominal quarterly return to unhedged US bond returns have therefore been reduced by 42 percent to reflect the impact of taxation and transaction costs. For the detail shown in the right-hand panel of Figure 4, the only assumption that has been altered relates to the estimated return to wine. Specifically, in the efficient frontier plots shown in the right-hand panel of Figure 4, it has been assumed that transaction costs and storage costs reduce the return to wine by 40 percent rather than 30 percent.

Let us first consider the left-hand side figure. For the sample period, using the above assumptions, unhedged US shares provided the highest quarterly after-tax and transaction cost return, but also had the highest risk. Australian bonds provided the lowest risk but also had the lowest return. To see the value of wine as an investment good, consider the two efficient investment frontier plots. The first frontier plot shows the minimum portfolio risk achievable for a given level of return when the only investments held are Australian shares, Australian bonds, US bonds, and US shares. The second frontier shows the minimum portfolio risk achievable when Australian wine is added to the portfolio mix. The extent to which the frontier shifts to the left represents the risk diversification gains due to wine.

In mean-variance analysis all points on the efficient frontier are by definition efficient. Yet for illustration purposes it is helpful to focus on a single portfolio. From the set of efficient portfolios, it is possible to argue that the portfolio providing the highest return per unit of risk

is the optimal portfolio. As this approach to identifying an optimal portfolio does not rely on utility theory, it is an intuitively attractive option. The Sharpe ratio measures return (or excess return) per unit of risk (standard deviation) and here has been calculated as (after-tax and costs portfolio return/portfolio risk). When wine is included in the portfolio mix, the optimal portfolio has quarterly return 1.81 percent, and risk 2.99 percent. To achieve this same level of portfolio return without including wine in the portfolio mix requires accepting portfolio risk of 3.23 percent per quarter. So, for the portfolio that maximises the Sharpe ratio, wine results in a gain in terms of risk reduction of 0.24 percent per quarter.

The right-hand panel of Figure 5 shows that, even if we assume transaction and storage costs reduce the return to wine by 40 percent, wine still provides portfolio risk reduction benefits. In this case, when wine is included in the portfolio mix, the optimal portfolio has return 1.77 percent, and risk 3.01 percent. Without wine in the portfolio mix, the risk associated with the portfolio giving quarterly return 1.77 percent is 3.18 percent. So, for the portfolio that maximises the Sharpe ratio, wine results in a gain in terms of risk reduction of 0.17 percent per quarter. The figure shows that even when using an unrealistically high figure for the effect of transaction and storage costs on the return to wine, adding wine to an already diverse investment portfolio results in a gain in terms of portfolio risk reduction.

FIGURE 5. THE RISK DIVERSIFICATION BENEFITS OF WINE INVESTMENT
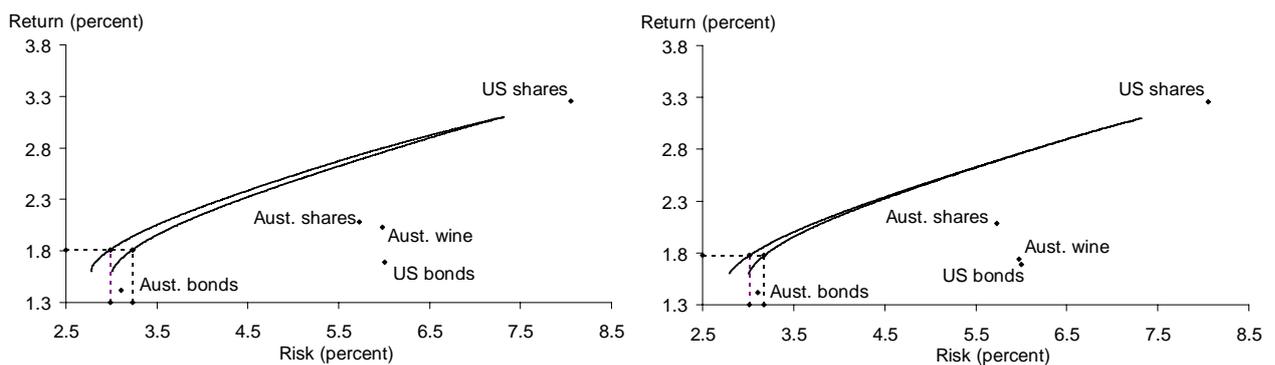
4. CONCLUSION

Ever since Krasker (1979) first posed the question of whether wine should be saved or savoured, the main theme of the return to wine literature has been to say wine should be savoured and not saved. This paper has argued that the true value of wine as an investment good for the UK and Australian markets has generally been understated because: (i) quality vintages can be identified in advance and investors hold only quality vintages; (ii) wine investment generally enjoys a special tax status in these countries; and (iii) including wine in an investment portfolio allows portfolio risk to be reduced. Wine should always be savoured, but those who choose to both save and savour wine may be making a wise decision afterall.

REFERENCES

Armitage, S. 2004. "Returns After Personal Tax on UK Equity and Gilts, 1919-1998." *European Journal of Finance,* 10: 23-43.

Ashenfelter, O., D. Ashmore and R. Lalonde. 1995. "Bordeaux Wine Vintage Quality and the Weather." *Chance*, 8(4): 7-14.

ASX (various) AXS Investment Sector Report.

Burton, J. and J.P. Jacobsen. 2001. "The Rate of Return on Wine Investment." *Economic Inquiry*. 39: 337-350.

Byron, R. P. and O. Ashenfelter. 1995. "Predicting the Quality of an Unborn Grange." *Economic Record*. 71(212): 40-53.

De Vittorio, A. and V. Ginsburgh. 1996. "Pricing Red Wines of Medoc Vintages from 1949 to 1989 at Christie's Auctions." *Journal de la Societe de Statistique de Paris*. 137: 19-49.

Gallagher D.R. and Jarnecic E. 2002 "The Performance of Active Australian Bond Funds.", *Australian Journal of Management*, 27: 163-185.

Gladstones, J. 1992. Viticulture and Environment. Winetitles, Adelaide, Australia.

Fogarty, J.J. 2006. "The Return to Australian Fine Wine." *European Review of Agricultural Economics*, 33: 542-561

Frino A. and D.R. Gallagher 2001. "Tracking S&P 500 Index Funds: Getting Your Money's Worth?" *Journal of Portfolio Management*, 28: 44-55.

Jorion, P. 1985. "International Portfolio Diversification with Estimation Risk." *Journal of Business*. 58: 259-278.

Krasker, W.S. 1979. "The Rate of Return to Storing Wine." *Journal of Political Economy*. 87: 1363-1367.

Ledoit, O. and M. Wolf. 2003. "Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection." *Journal of Empirical Finance*. 10: 603-621.

Markowitz, H. 1952. "Portfolio Selection." *Journal of Finance*. 7: 77-91.

Sharpe, W.J. 1991. "The Arithmetic of Active Management." *Financial Analysts Journal*. 47: 7-9.